



Grid Computing

# Grid Computing Basics

**Nancy Brittle**, Grid Computing Sales Exec, Americas  
email: [nbrittle@us.ibm.com](mailto:nbrittle@us.ibm.com)

**Michael Osias**, IT Grid Architect  
**Duane Quintern**, IT Grid Specialist



April 27, 2004

# Agenda

## § Grid Basics

### **What is a grid?**

Definition

Uses & Types

Examples

### **How and why grids are deployed**

### **Building a grid**

Considerations

Tools

### **Resources and 'Getting Started'**

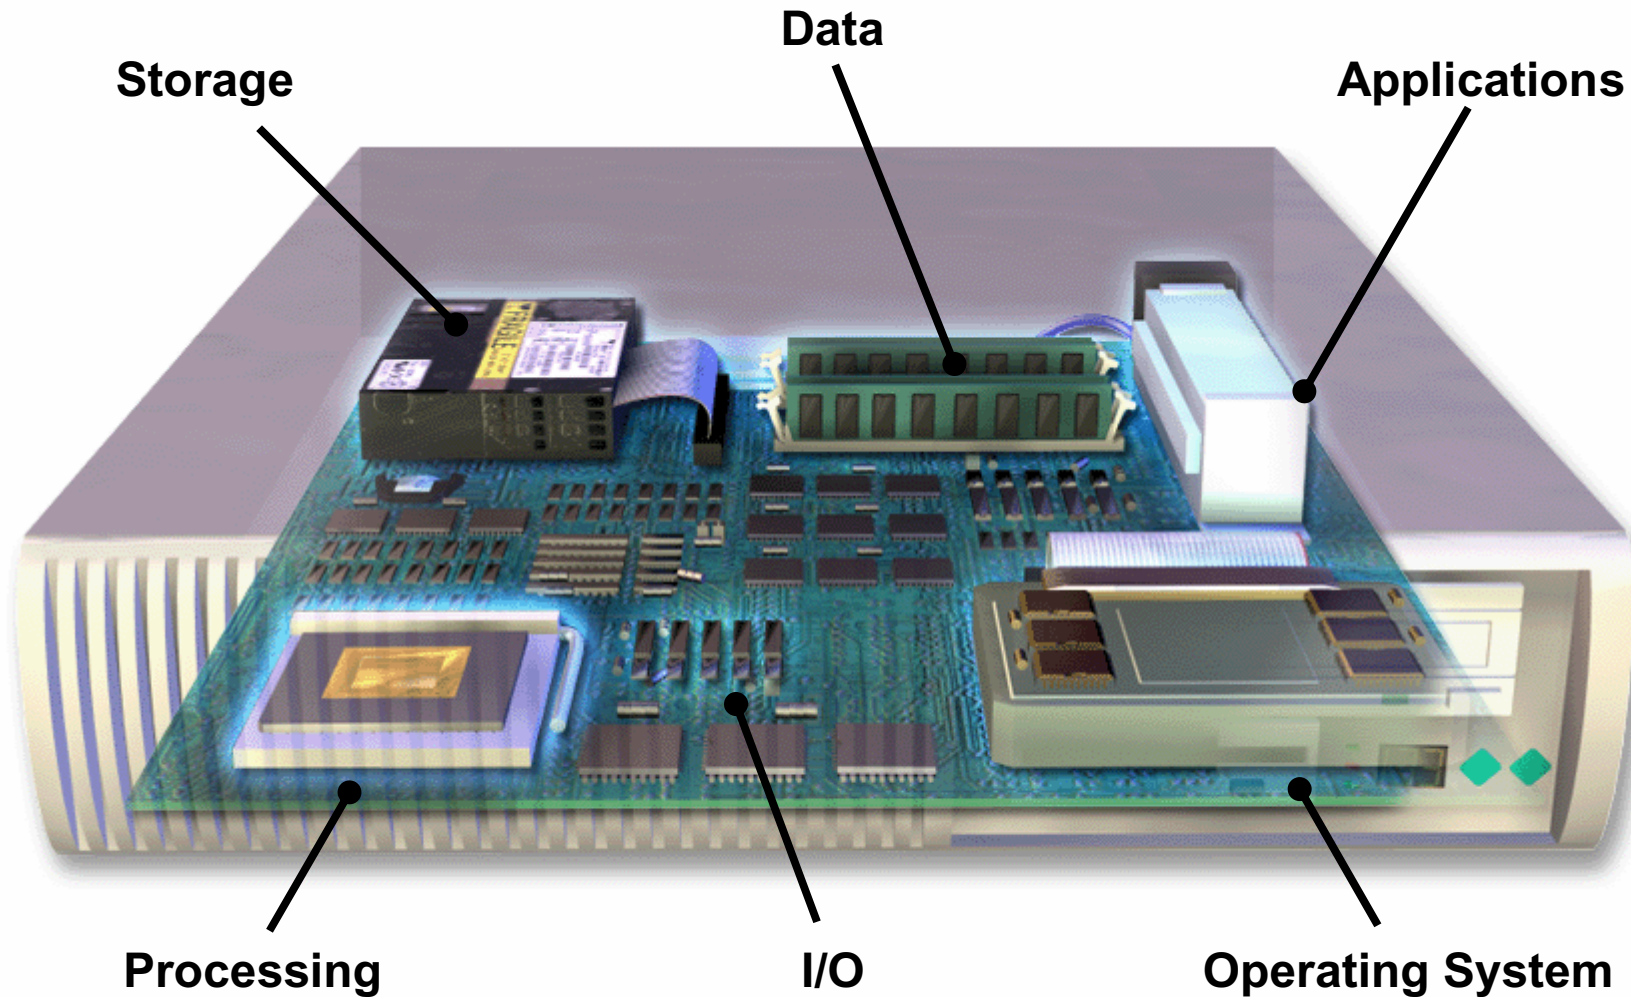
# What is a grid?

Definitions, Uses, Types, and Examples



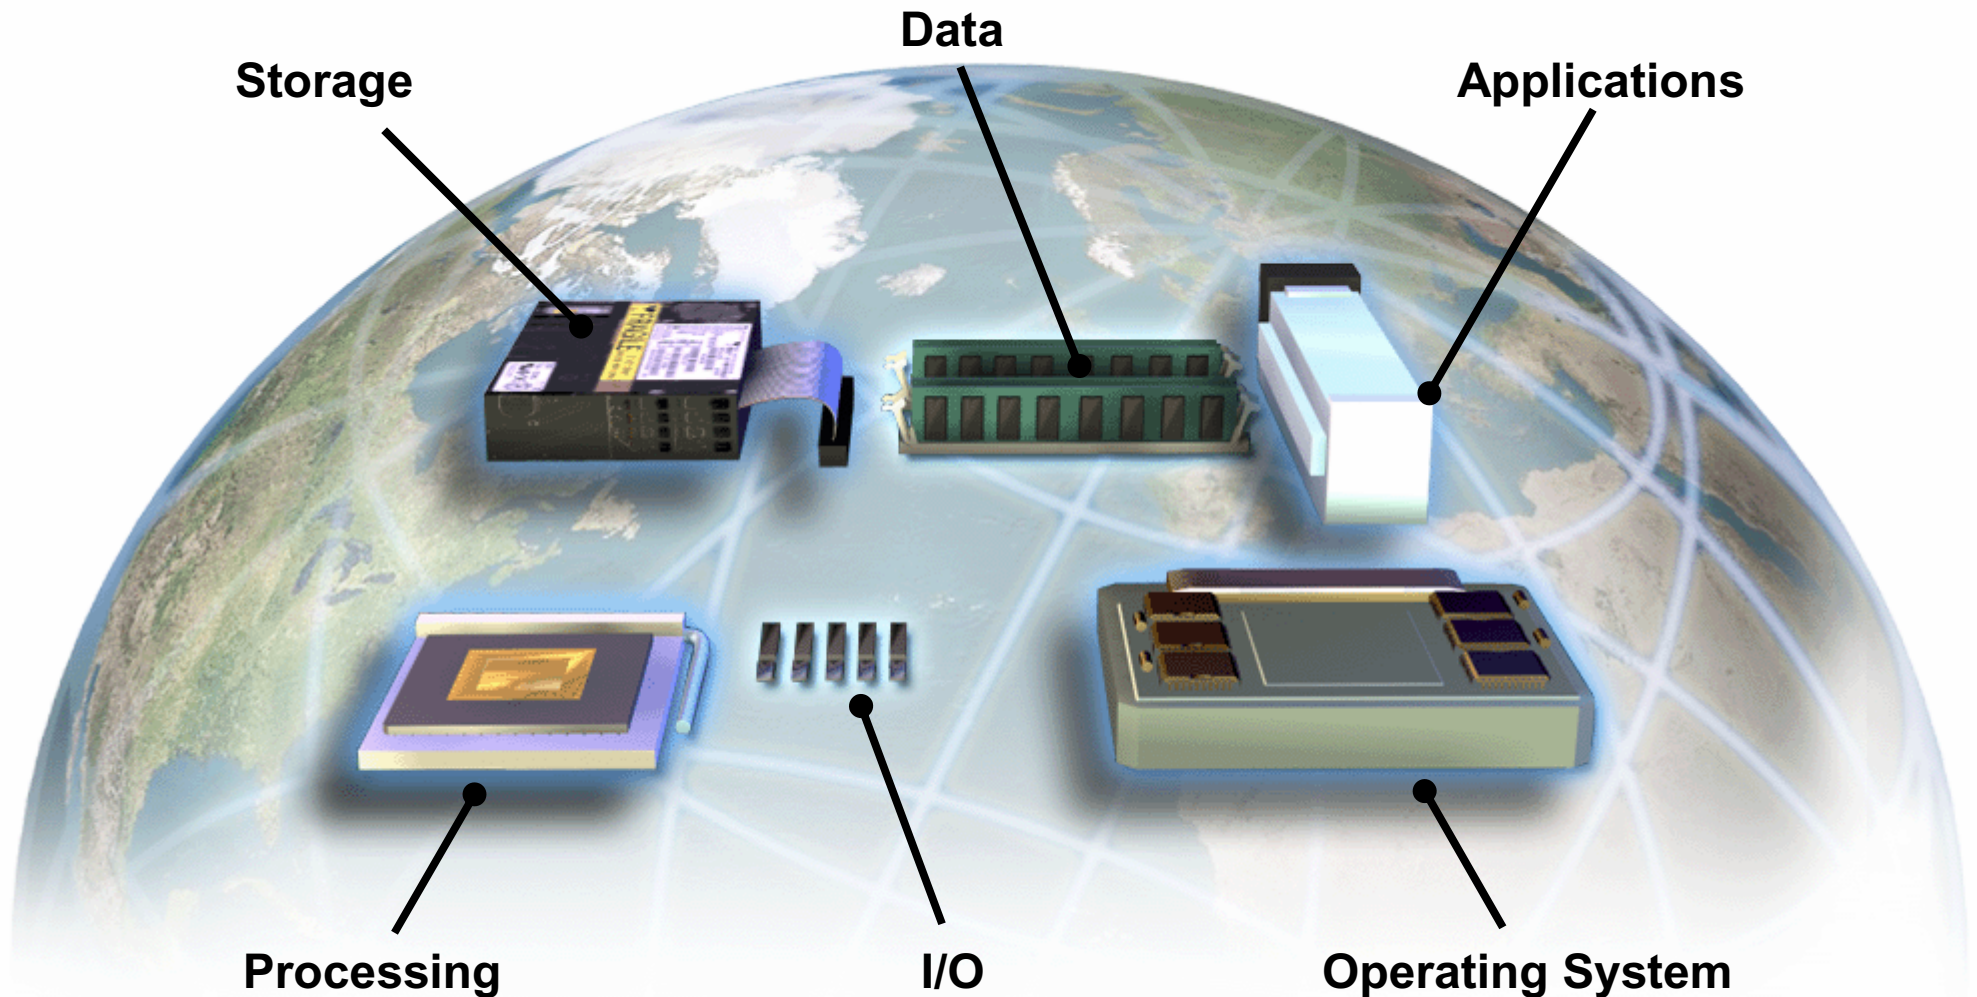
# Grid Computing

## *Microcosm – Pre-Internet “System”*



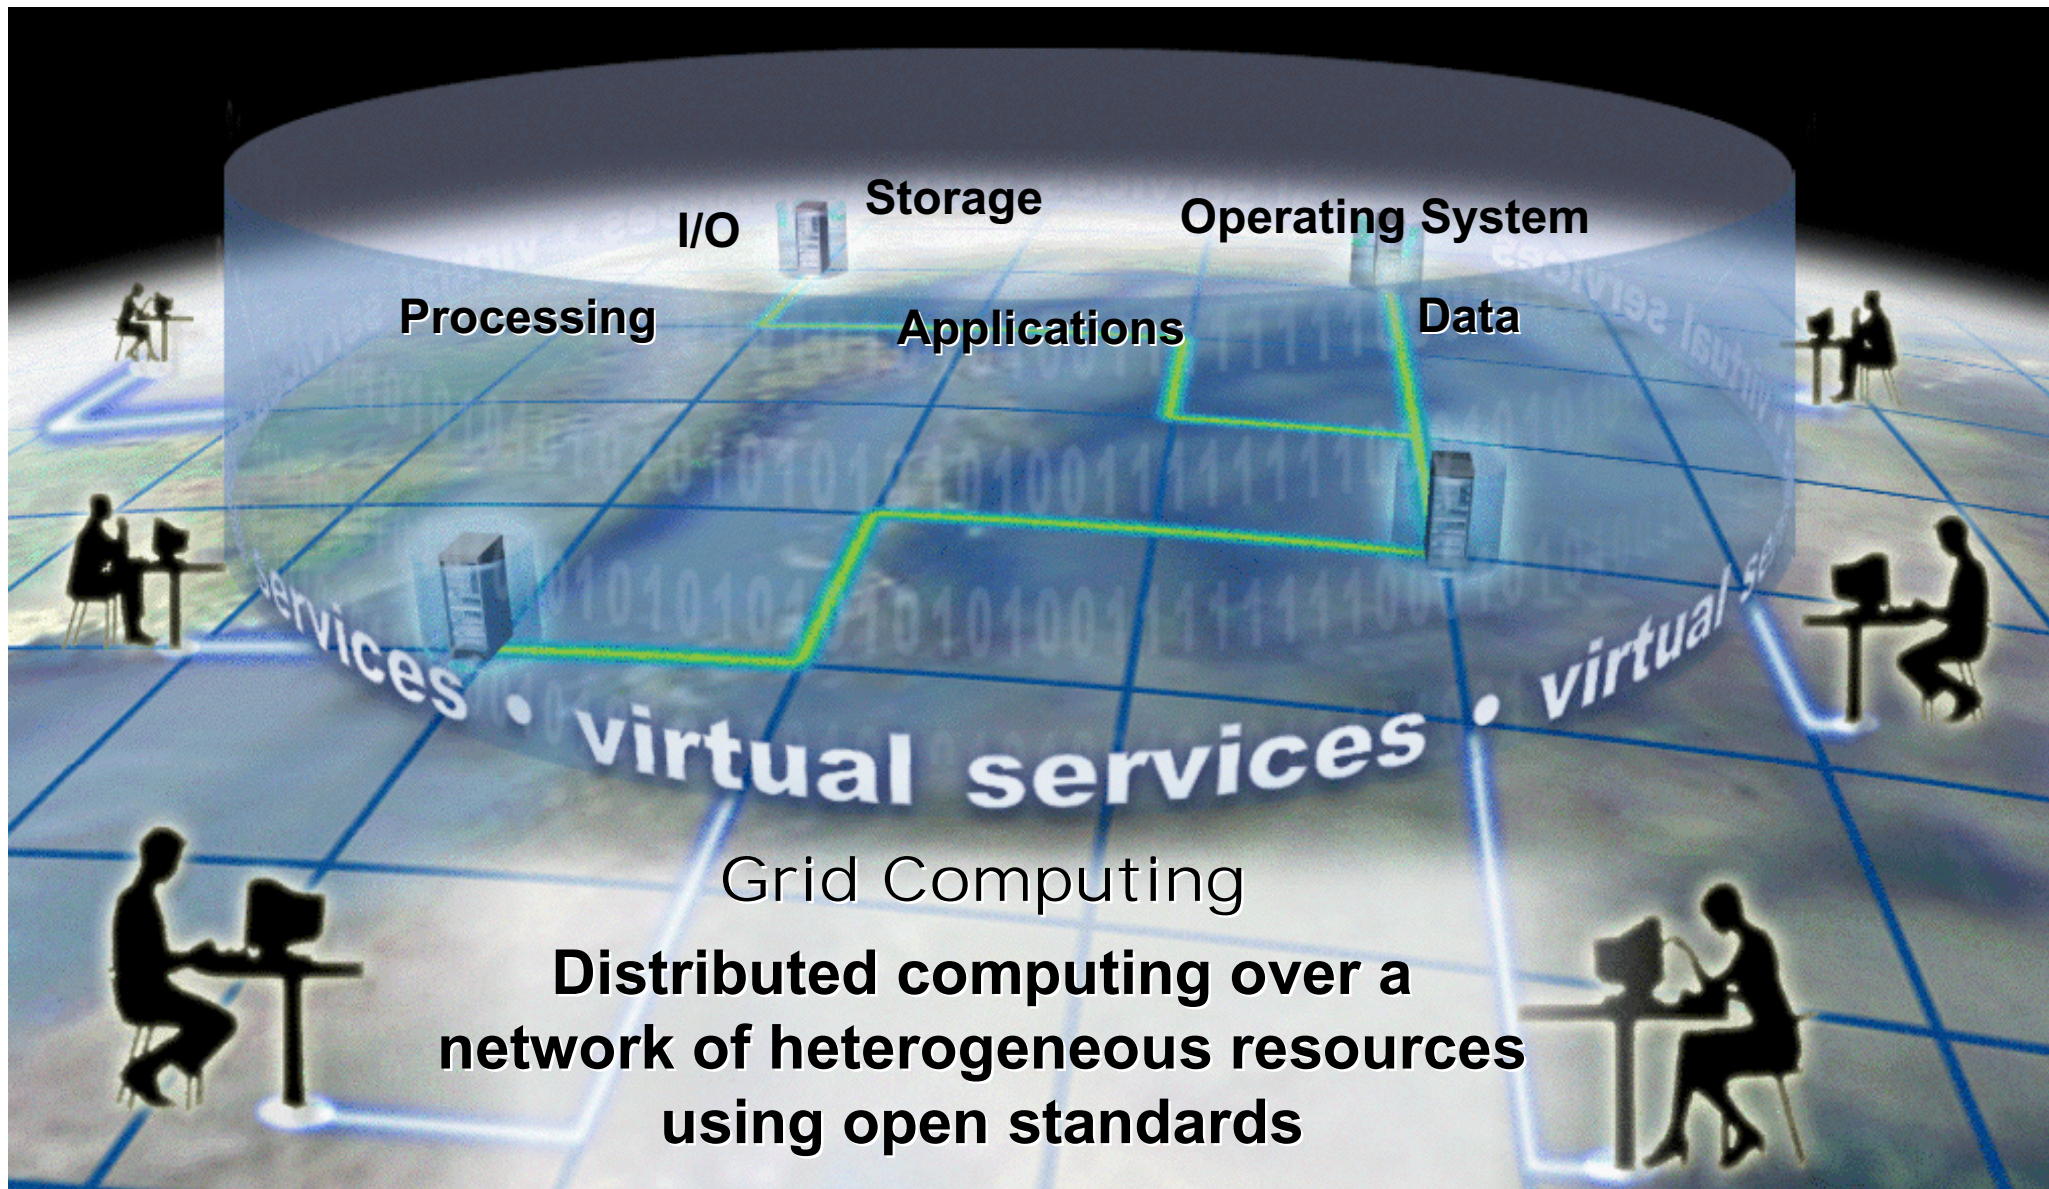
# Grid Computing

## *Macrocosm – Distributed Resources and Applications*



*....a single unified image*

# Grid Computing



# **Grid Computing Resources & Types**

## **Grid Resources**

- § **Computation**
- § **Storage**
- § **Data**
- § **Applications**
- § **Communication (I/O)**
- § **Software & Licenses**
- § **Special equipment, capacities, architectures, & policies**

## **Grid Types**

- § **Collaboration Grid**
- § **Compute Grids**
  - Desktop Scavenging
  - Server
- § **Data/Information Grids**
  - Content
  - Data
  - File
  - Storage



**Grid Resources Virtualized Across the Grid Types**



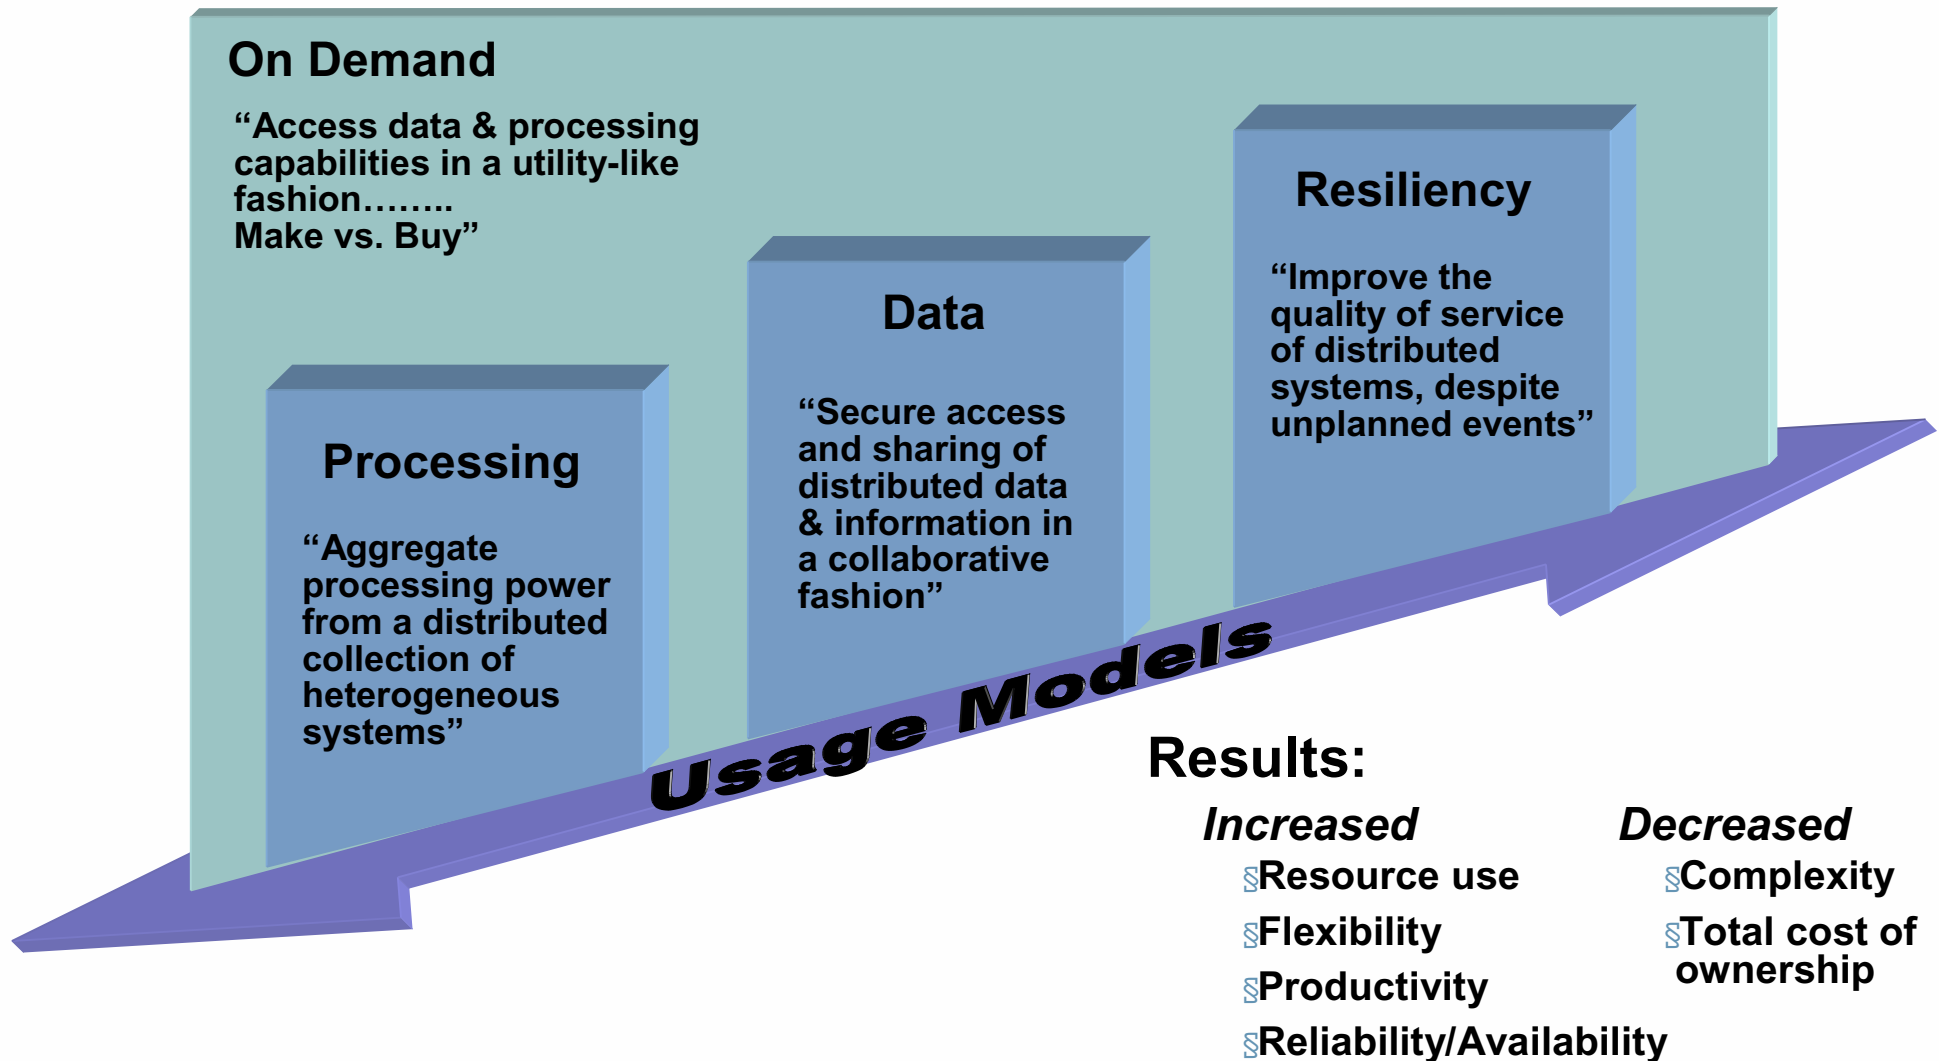
Grid Computing

# How and Why Grids are Deployed

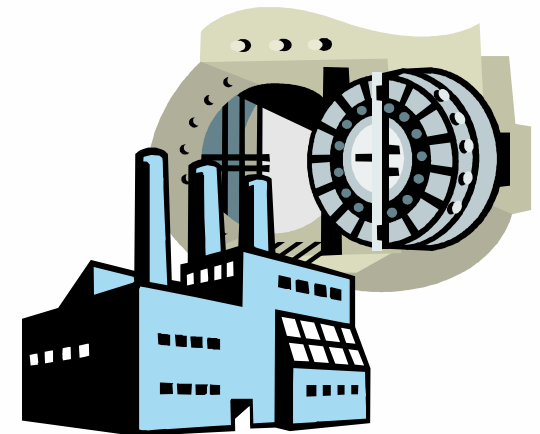
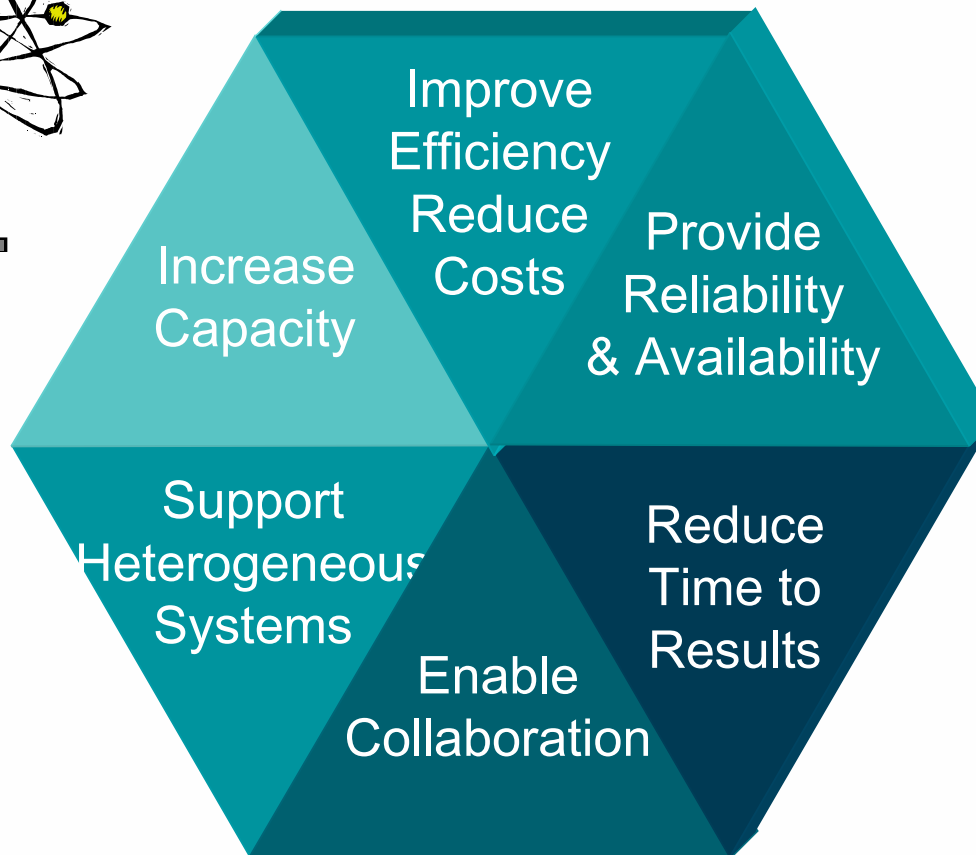
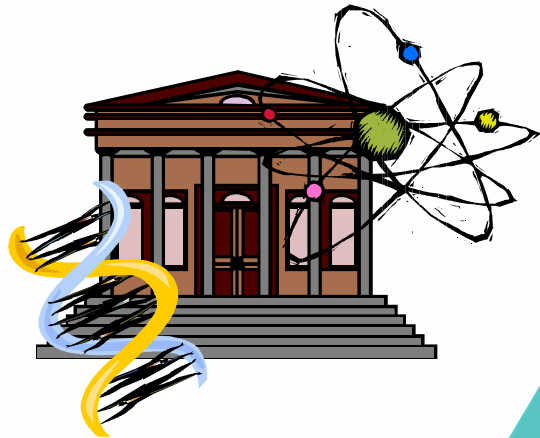


# Grid Computing Value Proposition

## 3 Models and Unique Value Propositions



# Motivations for Grid Computing

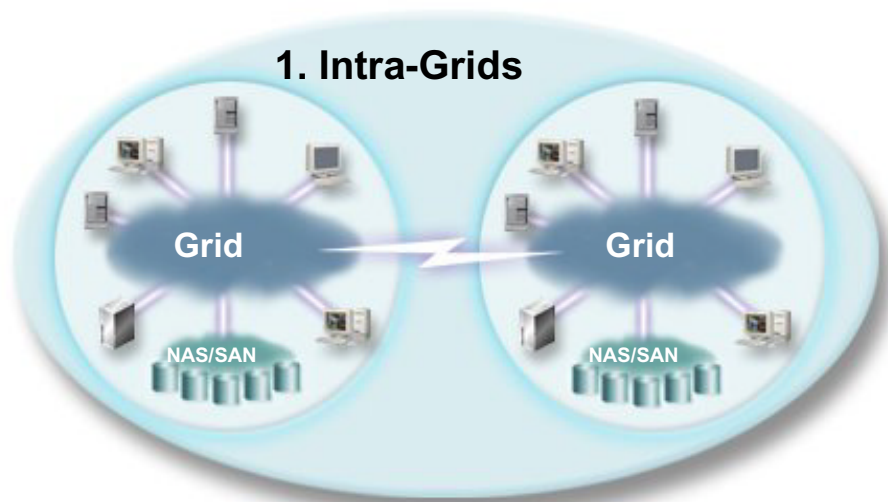


# Server / Storage Utilization

	Peak-hour Utilization	Prime-shift Utilization	24-hour Period Utilization
<b>Mainframes</b>	<b>85-100%</b>	<b>70%</b>	<b>60%</b>
<b>UNIX</b>	<b>50-70%</b>	<b>10-15%</b>	<b>&lt;10%</b>
<b>Intel-based</b>	<b>30%</b>	<b>5-10%</b>	<b>2-5%</b>
<b>Storage</b>	<b>N/A</b>	<b>N/A</b>	<b>52%</b>

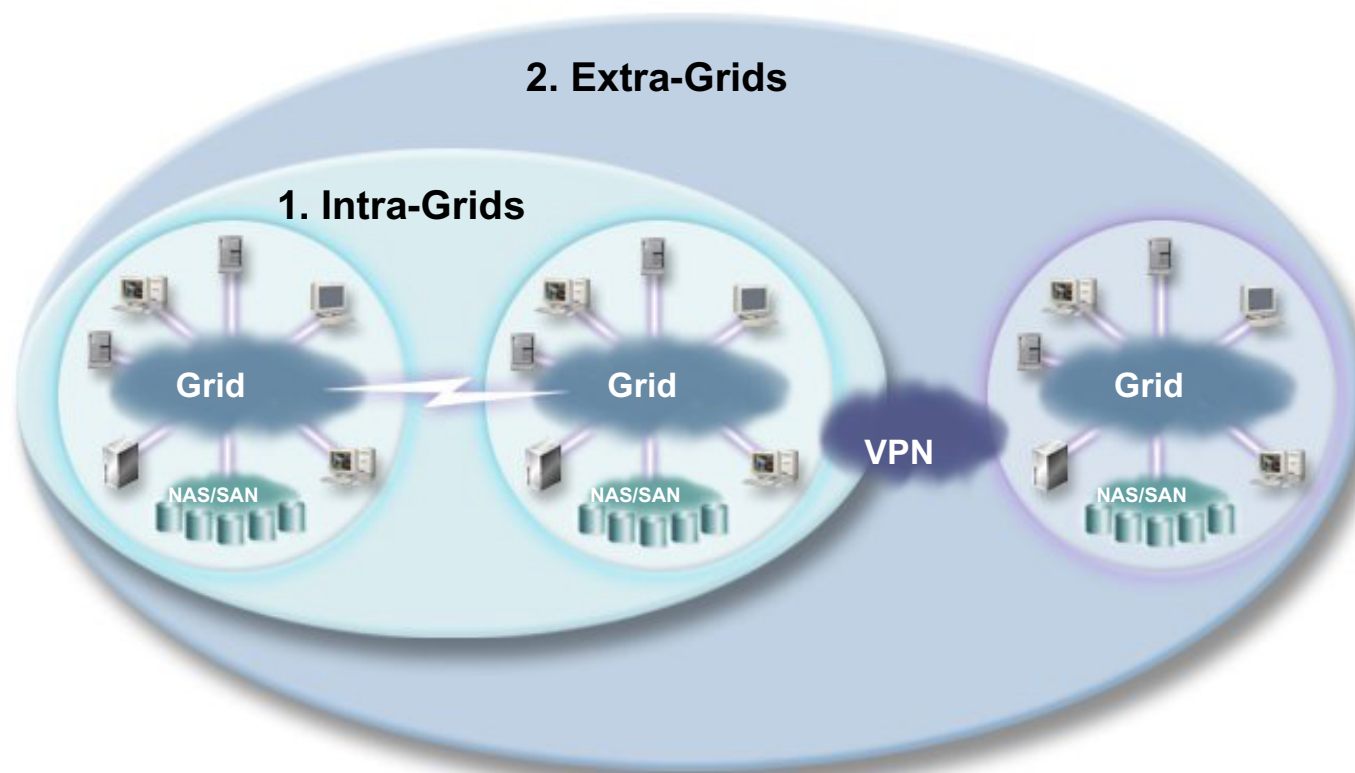
# Grid Deployment Options

*A Function of Business Need, Technology and Organizational Flexibility*



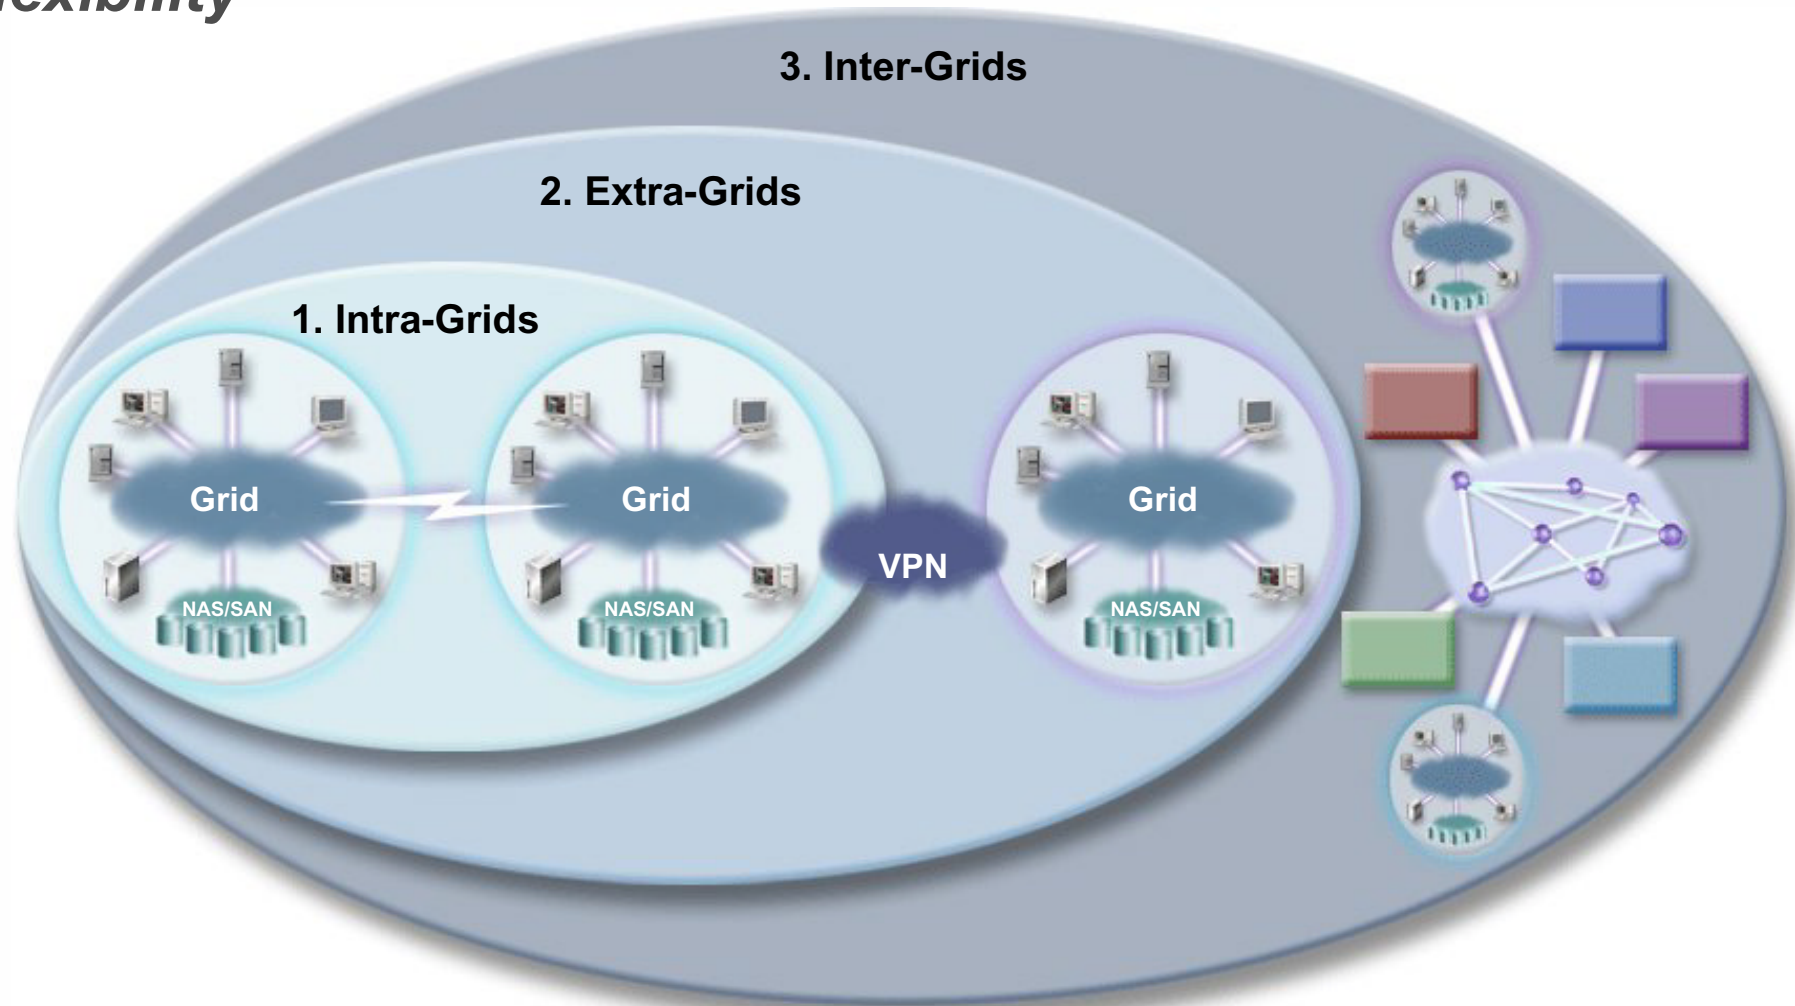
# Grid Deployment Options

*A Function of Business Need, Technology and Organizational Flexibility*



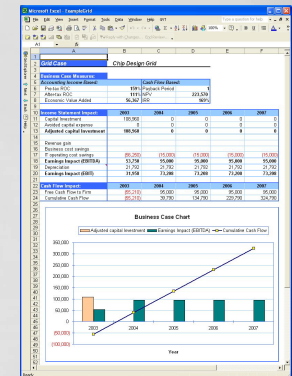
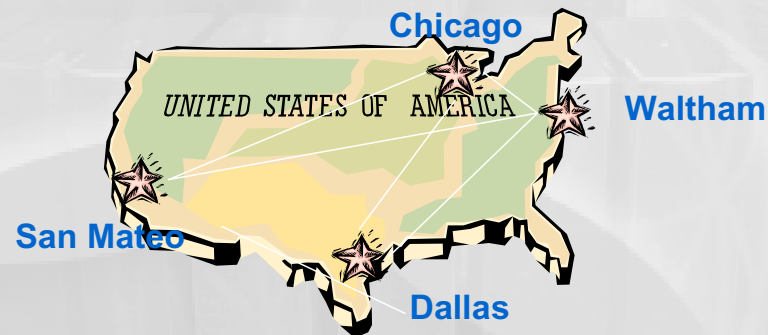
# Grid Deployment Options

*A Function of Business Need, Technology and Organizational Flexibility*

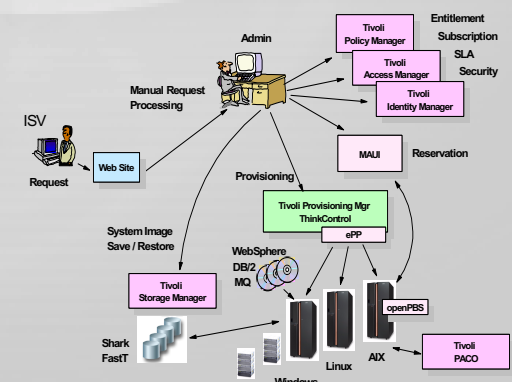




## PC World



## Solutions Grid for Business Partners



## Virtual Loaner Program



## The TeraGrid – Extensible Terascale Facility (ETF)

### National Science Foundation Grid Computing project (\$90M):

- § Connects nine (9) major supercomputing sites: NCSA, SDSC, Argonne NL, CalTech, PSC, UTexas, IndianaU, PurdueU, Oak Ridge NL

- § 40 gigabit network backbone connecting the sites
- § 20 Teraflops of computing power
- § 1 Petabyte of disk accessible data storage
- § Accessible to thousands of scientists working on advanced research

#### § Applications include:

- § Real Time Brain Mapping
- § Earthquake Modeling
- § Molecular Dynamics simulation
- § Mcell – Monte Carlo simulation of cellular micro physiology
- § Encyclopedia of Life – Protein catalog

#### § IBM project team and solution includes:

- § IBM High Performance Computing (HPC) expertise
- § IBM GPFS expertise
- § IBM Linux Clusters – Itanium2 processors
- § IBM Power4 processors – p690 Regattas
- § IBM Grid Computing & Linux consulting services



# DEISA

Distributed  
European  
Infrastructure for  
Supercomputing  
Applications

## § A European funded project to create a Virtual Supercomputer based on existing National Supercomputing Centers

- IDRIS in France
- CINECA in Italy
- IPP, Juelich in Germany
- SARA in Netherlands



## § IBM project team and solution includes:

- IBM Power4 processors – p690 Regattas
- IBM High Performance Computing (HPC) expertise
- IBM GPFS expertise
- IBM Grid Computing expertise

# Deutscher Wetterdienst (DWD)

*Research & Development*

## German Meteorological Office

### Challenge:


DWD wanted to make available its significant compute resource available to the Federal Authorities for Hydraulic Engineering (BAW) and Maritime Navigation (BSH) in order to better undertake numerous calculation-intensive simulations of the varying water environment and better understand the risk of flooding

### Solution:

Incorporate leading edge hardware and middleware to establish a grid linking DWD, BAW and BSH.

The solution includes:

- § IBM RS/6000 Unix server with 2000 processors
- § Unicore grid middleware



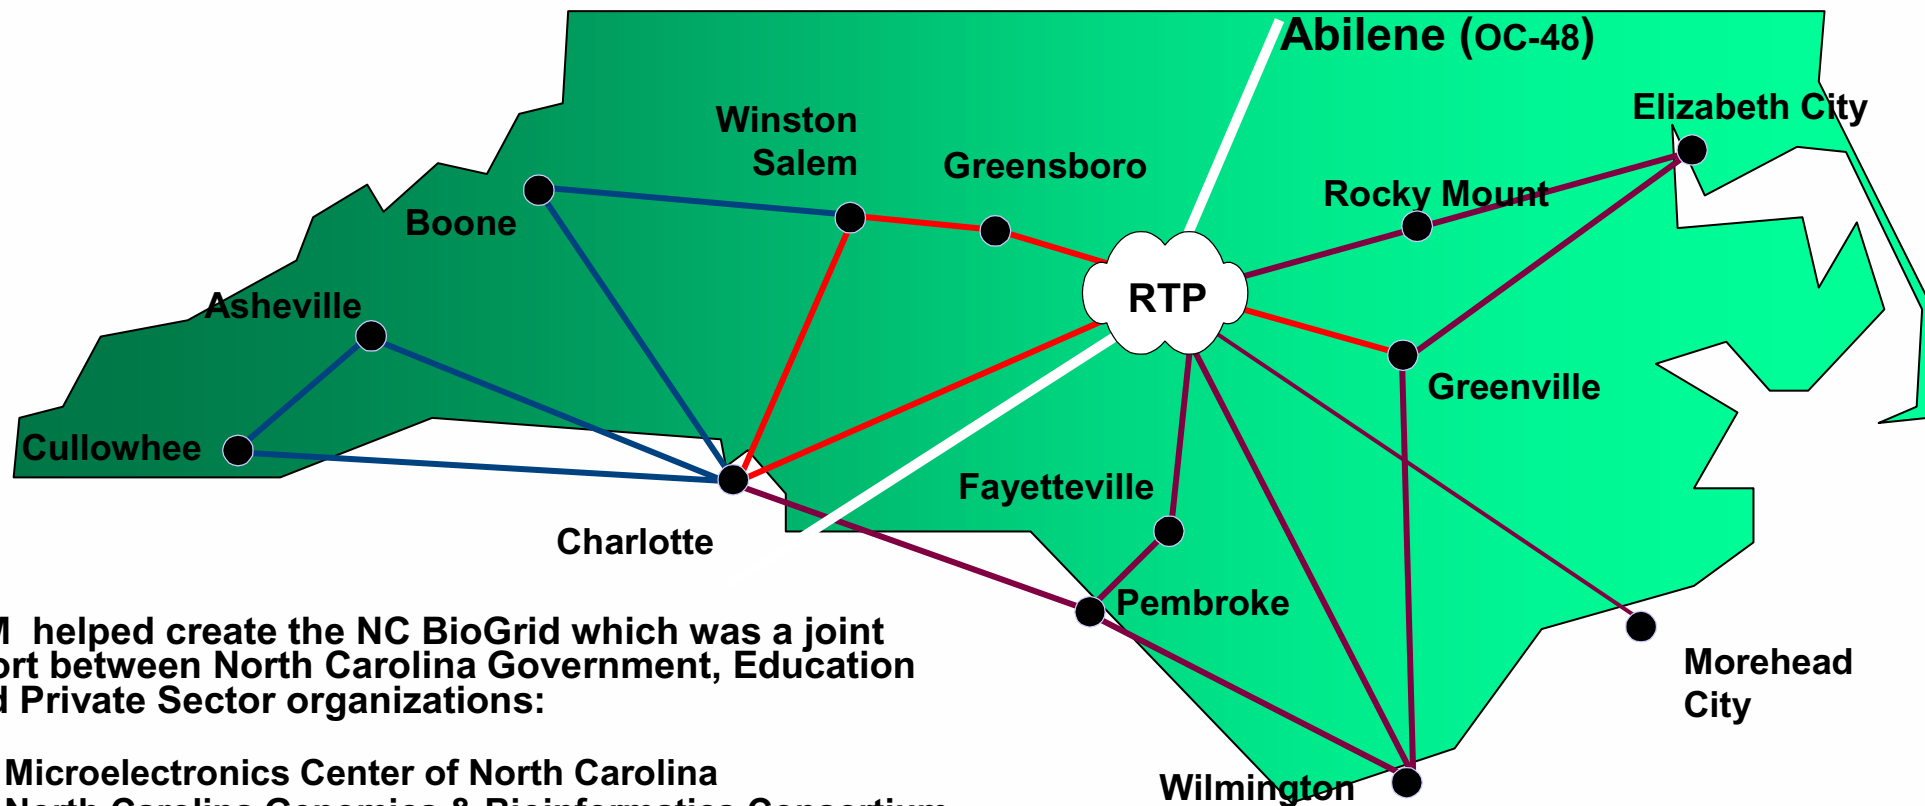
Deutscher Wetterdienst

**Benefits:**

- § Flexible sharing of IT resources between the partners
- § Less queuing time for jobs resulting in reduced waiting time and faster execution of processes
- § Potential to expand the scope of the project to include other authorities

**“The computer network between DWD, BAW and BSE highlights the strong benefits especially public authorities can gain from Grid technology. With the help of the Unicore software, complicated authorisation processes are no longer necessary. The users now can place their jobs as easy as before from wherever they are.” Geerd-Ruediger Hoffmann, Head of Systems, DWD**

## North Carolina BioGrid



IBM helped create the NC BioGrid which was a joint effort between North Carolina Government, Education and Private Sector organizations:

- Microelectronics Center of North Carolina
- North Carolina Genomics & Bioinformatics Consortium
- North Carolina Supercomputing Center
- North Carolina Research and Education Network
- University of North Carolina
- Duke University
- North Carolina State University
- GlaxoSmithKline & BioGen
- SAS

# NDMA: National Digital Mammographic Archive

- § Combines Grid Computing with Radiology to make breast cancer diagnosis faster and treatment more effective
- § IBM assisted with implementing a Grid infrastructure across the hospitals to manage and retrieve digital mammograms (256mb each)
- § Secure transmission of all patient records
- § Grid solution architecture includes: IBM pSeries, xSeries, Linux, DB2, GPFS
- § WebSite: <http://nscp.upenn.edu/NDMA>  
<http://www.i3archive.com/>

## Research & Development



THE DIVISION OF THE BIOLOGICAL SCIENCES



# China Grid

**Launched by China Ministry of Education (MOE) in August 2003**

**Phase I planned to be completed in 2005:**

Cover 12 universities across China  
6 Tflops and over 60TB storage capacity

**Will cover 100 universities in the near future**

**Focus on 5 applications:**

Life Science  
e-learning  
Mass Information Processing (Digital Olympic)  
Data Grid  
Fluid Dynamics

**IBM provided solution & tech. support including:**

eServer pSeries & xSeries  
Storage system  
IBM Global Services  
IBM Research

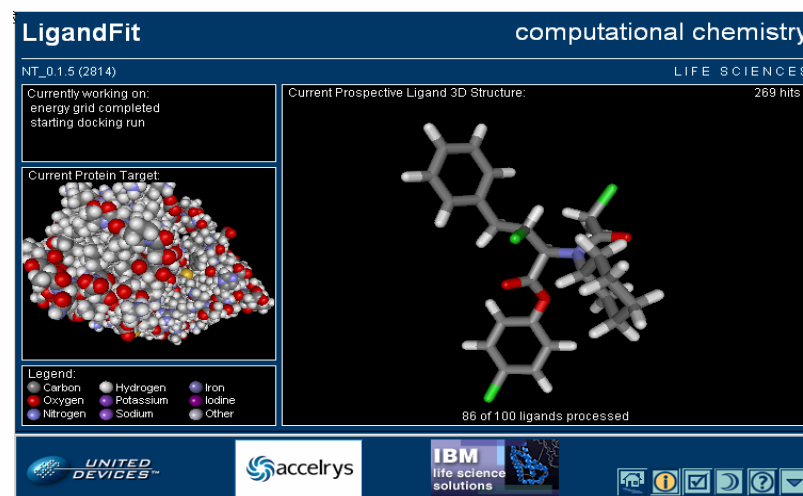
## Research & Development



# The SMALLPOX Research Grid

- § A massive distributed computing grid running a computational chemistry application to help fight the smallpox virus
- § The grid is screening 35 million potential drug molecules against several protein targets. Two million computing devices are connected to this grid
- § The Grid architecture will reduce the time required to develop a commercial drug by several years
- § IBM is collaborated with United Devices, Accelrys, Evotec OAI, US Department of Defense and Oxford University on this project
- § IBM is providing the infrastructure technology, high-performance computing and data management software (p690, AIX, DB2, Linux)

## Research & Development





Grid Computing

# Building a grid

Considerations, Tools



# Cooperation on Standards



# The Value of Open Standards

Distributed Computing:

*Grid Services*

*(OGSA, WS Resource Framework,  
GGF -> Globus)*

Applications:

*Web Services*

*(SOAP, WSDL, UDDI)*

Operating System:

*Linux*



Information:

*World-wide Web*

*(html, http, j2ee, xml)*

Communications:

*e-mail*

*(pop3, SMTP, Mime)*

Networking:

*The Internet*

*(TCP/IP)*

**e-business**

# Open Grid Services Architecture (OGSA)

## § Objectives:

- Manage resources across distributed heterogeneous platforms
- Deliver seamless QoS
- Provide a common base for autonomic management solutions
- Define open, published interfaces

## § Exploit industry-standard integration technologies

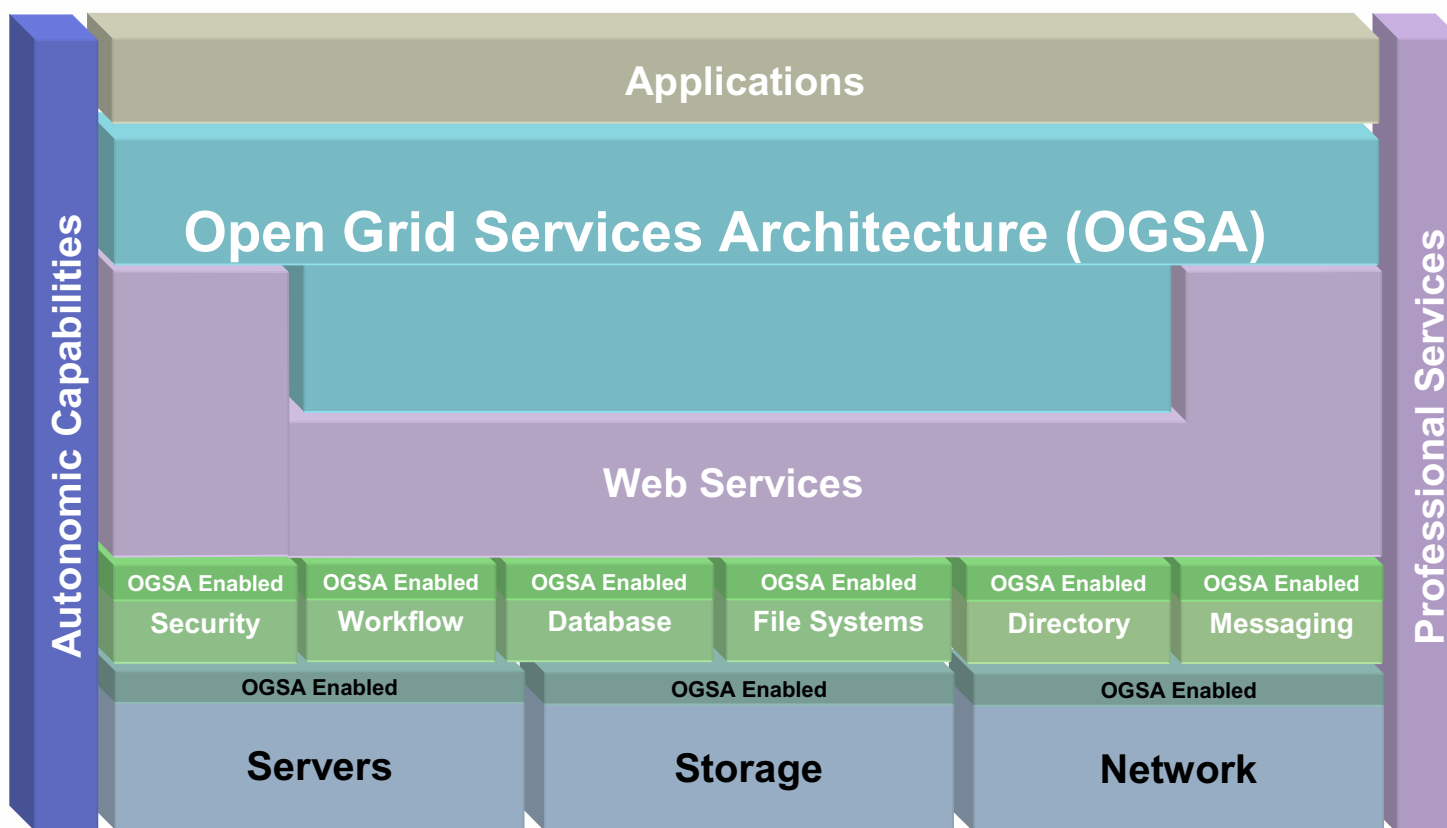
- Web Services: SOAP, XML, WSDL, WS-Security, UDDI...

## § Integrate with existing IT resources

# Open Grid Services Architecture (OGSA)

- § OGSA introduces extensions to web services
  - Autonomic capabilities
  - Provisioning
  - State management for web services
    - § Reliable, secure management of distributed state
- § Lines between applications and resources are blurring
  - OGSA extends key middleware as meta-os services
    - § Connectors and adapters virtualize native resources
    - § Logical rendering of physical resources
      - § Service compensation layer
    - § Resources rendered as services
  - Separate interface from access and implementation
- § OGSA provides Grid Programming Model Standardization

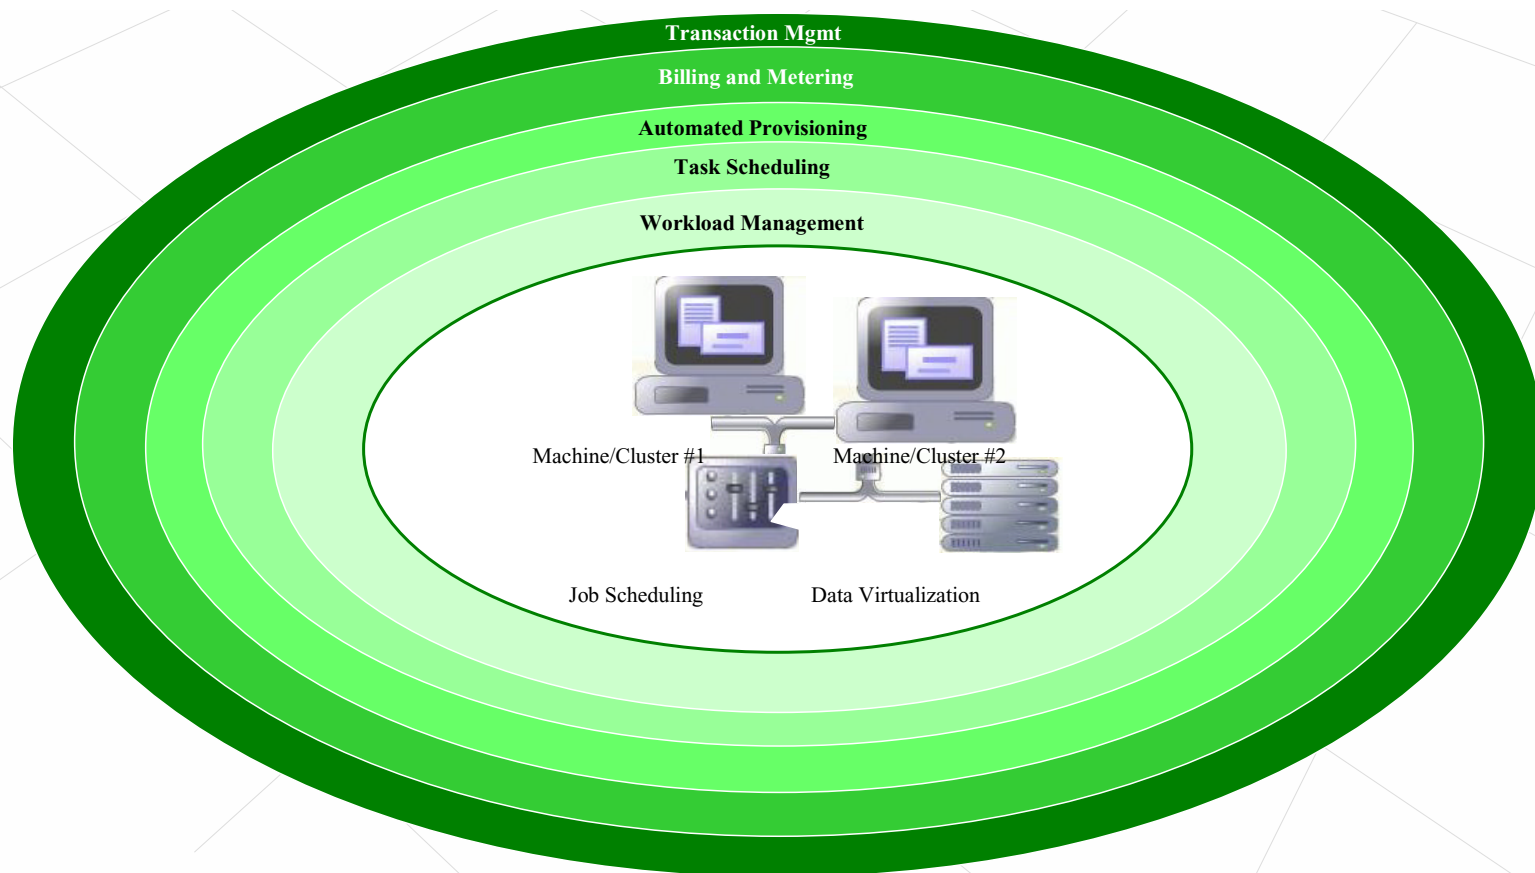
## Grid Solution Architecture Framework



# Grid Adoption Steps

## *Transaction Management:*

- Manage the execution of e-business transactions across distributed resources
- Enable dynamic allocation of resources for transactional and parallel application models



# On Demand Roadmap

On Demand Capabilities

Transaction Management

WebSphere.

Billing and Metering

Tivoli.

Workload Management

@server

Tivoli.

Information Virtualization

@server

TotalStorage™

DB2.

Tivoli.

AVAKI

Orchestrated Provisioning

@server

Tivoli.

Scheduling

WebSphere.

Tivoli.

DataSynapse

UNITED  
DEVICES™

Platform

Web Services/OGSA

@server

the photos alliance  
www.photosalliance.com

WebSphere.

# Information Virtualization

## **“Core component of the Grid Computing model...**

- § Allows end-users and applications secure access to any information source regardless of where it exists
- § Provides access to heterogeneous files, databases, or storage systems
- § Supports sharing of data for processing and/or large-scale collaboration

**...in an environment that is autonomic, virtualized, and open.”**

# Managing Information at Different Levels

## Data

- §Global Naming
- §Meta-data and catalog
- §Federation and Transformation

## File

- §Distributed File Systems / Remote Access
- §File Transfer / Data Replication
- §Caching

## Storage

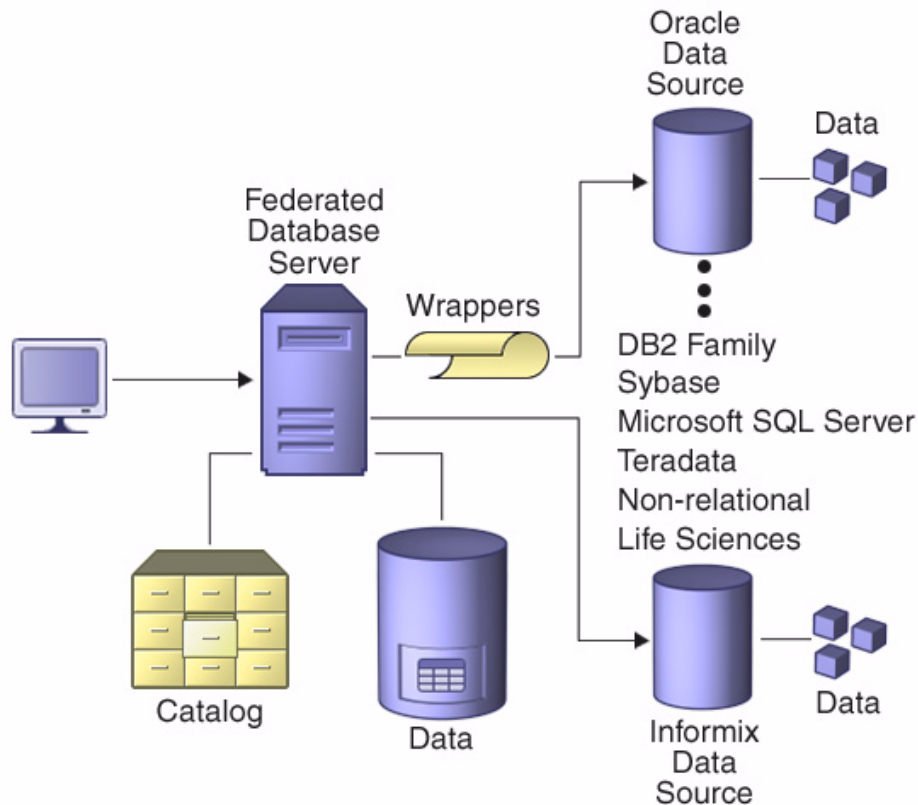
- §NAS / SAN “Storage Cluster”
- §Automatic or Dynamic provisioning of storage
- §Support for hierarchy management

# IBM Products for an Information Grid

Data	Product	Features	Benefits
	<u>DB2 Information Integrator</u>	Federated data server, replication server	Query and access distributed data without requiring central repository. Supports movement of data from mixed relational data sources.
	<u>DB2 UDB</u>	Relational database that runs on Linux, Unix, Windows, z/OS, and OS/390	Manageability features, Integrated Information capabilities via Web Services, Integrated business intelligence, and more
	<u>Avaki Data Grid 5.0*</u>	Data catalog, data provisioning, reusable data integrations, caching capabilities.	Provisioning, access, and integration of data from multiple, heterogeneous, distributed sources.
File	GPFS (General Parallel File System)	Cluster based, shared disk, parallel file system. Data and metadata can flow to all nodes and all disks in parallel. Featured in HPC environments. Available on pSeries and Linux clusters.	Not a client-server file system like NFS, DFS, or AFS: no single server bottleneck, no protocol overhead for data transfer.
	NFS v4	Provides scalable access to GPFS from outside cluster. GPFS + NFSv4 provides the performance of a SAN File System scalable to a WAN.	Security and access control in a grid environment.
	<u>SAN File System</u>	Provides a common file system specifically designed for storage networks. Manages the metadata on the storage network instead of within individual network servers.	Provides high performance access to data and enables sharing across heterogeneous application servers. Allows applications on any server within the SAN to access any file in the network without making changes to the application.
Storage	<u>SAN Volume Controller</u>	Creates pools of managed disks spanning multiple storage subsystems. Includes dynamic data-migration function.	Centralized point of control for volume mgmt. Allows administrators to migrate storage from one device to another w/o taking it offline.
	<u>Tivoli Storage Resource Manager</u>	Enterprise wide reporting, file level analysis, subsystem reporting, automated capacity provisioning	Storage on demand for file systems. Reclaim wasted space consumed by non-essential files. Ensure storage used efficiently for future capacity.
	<u>Tivoli Storage Manager</u>	Data backup/restore, data archive and retrieve	Centralized protection leading to faster backups and restores with less resources needed.

One type of Grid data virtualization is a virtual and consistent interface to heterogeneous federated data.

## ***Federation of Data***



### § Transparency

- Appears to be one source

### § Heterogeneity

- Integrates data from diverse sources
- Relational, Structured, XML, messages, Web, ...

### § Extensibility

- Federate almost any data source.
- Development tooling provided

### § High Function

- Full query support against all data
- Capabilities of sources as well

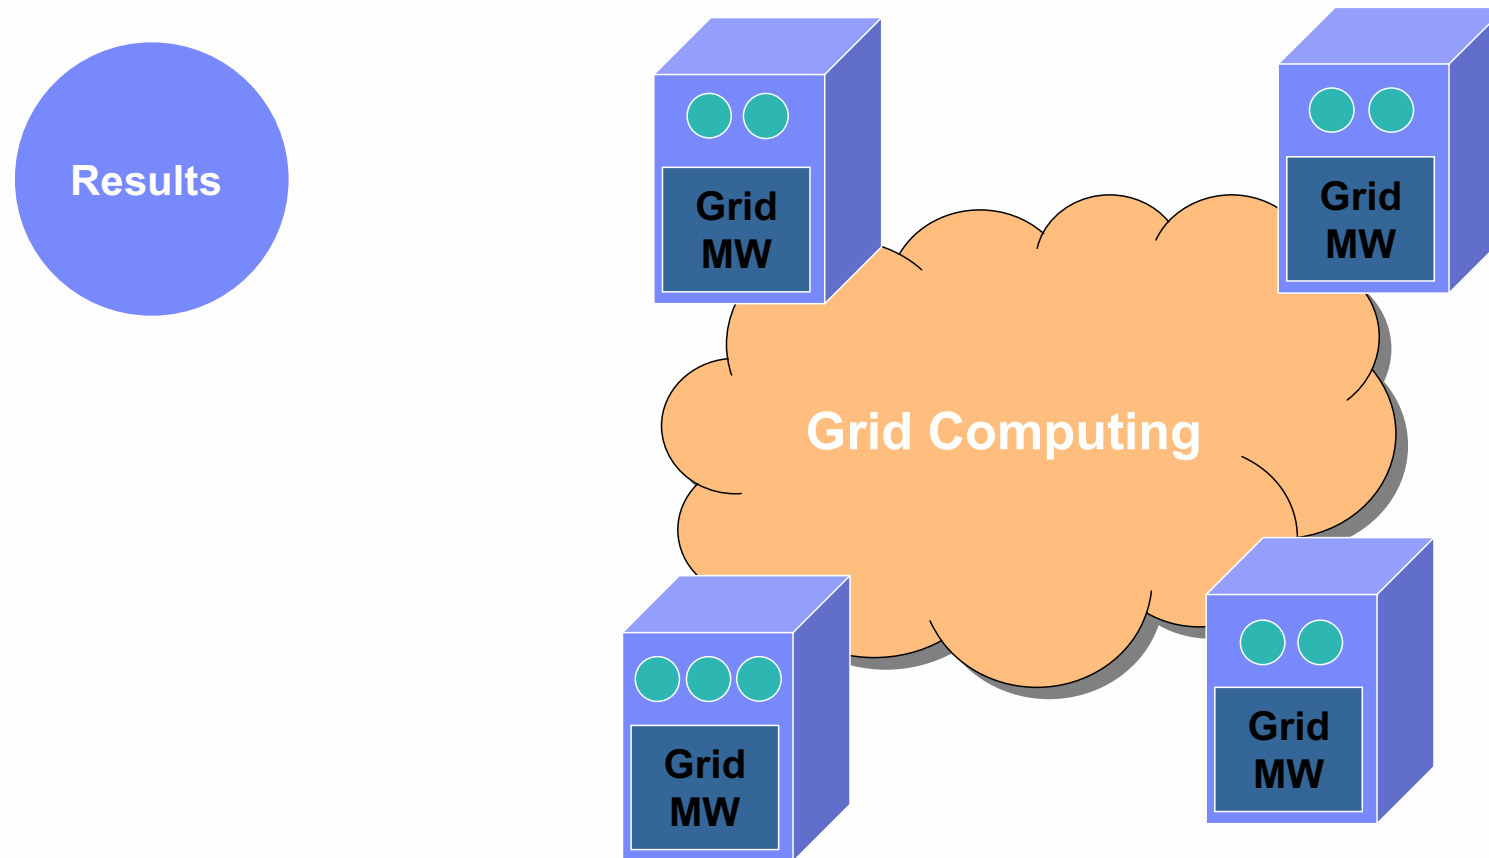
### § Autonomy

- Non-disruptive to data sources, existing applications, systems.

### § Performance

- Optimization of distributed queries

# Scheduling



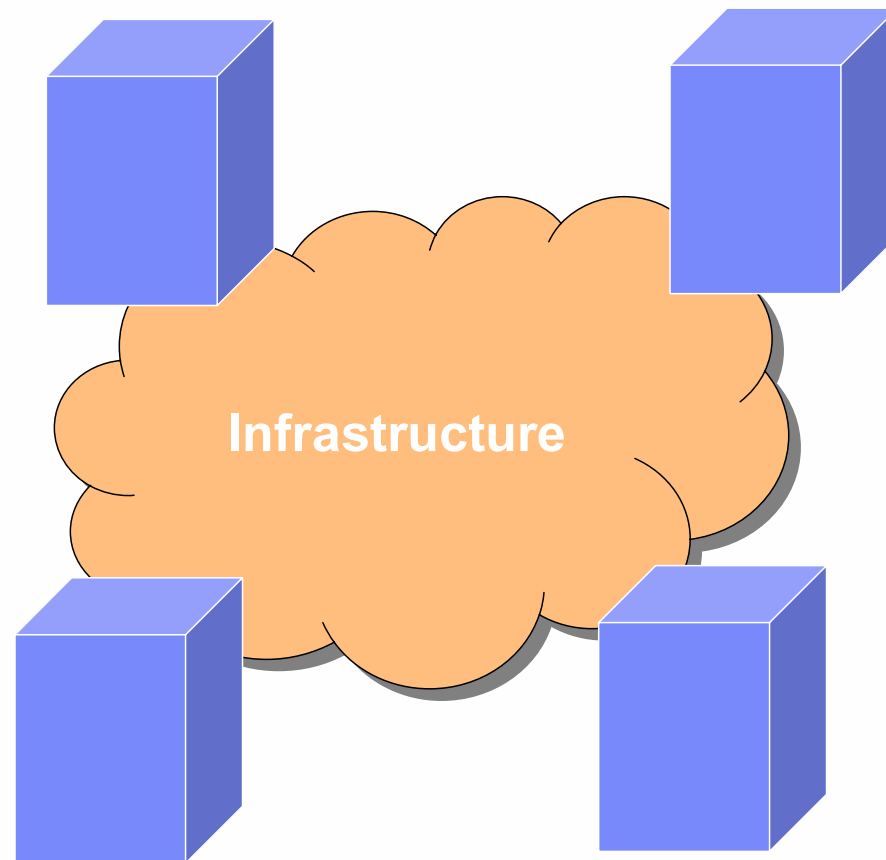
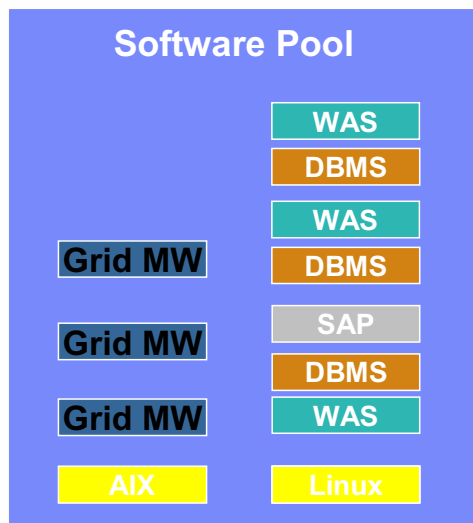
# Orchestrated Provisioning

## Need 3 Linux Servers

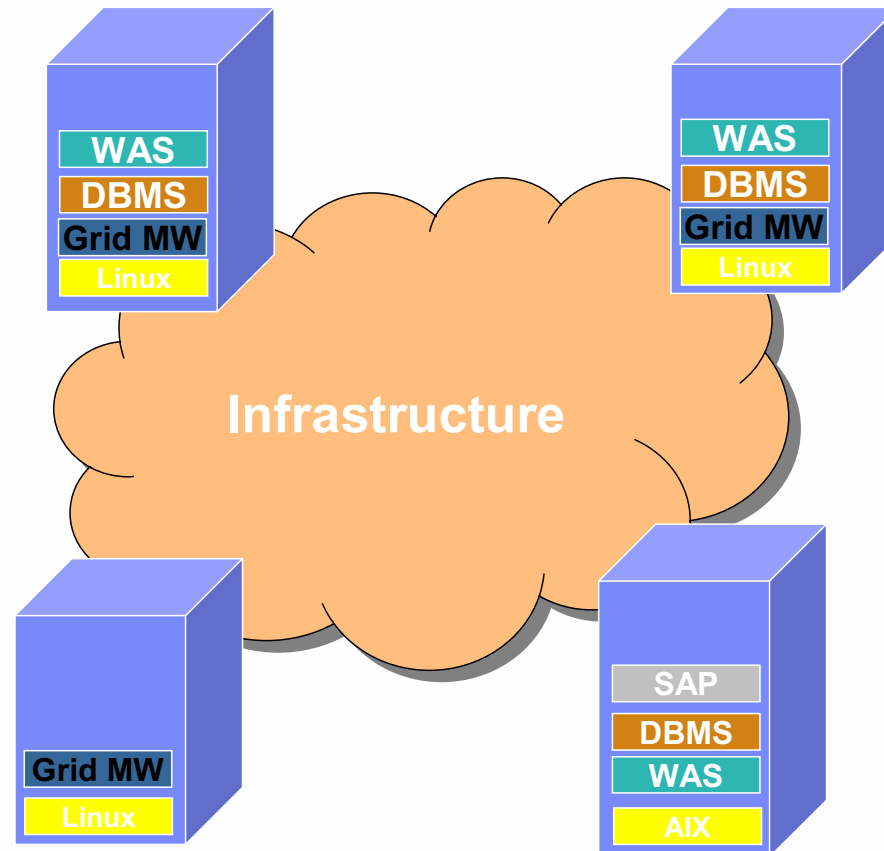
- § Intel, 2 Way Processors, 1 Gig of Memory, 20 Gig of Storage
- § Linux
- § Grid MW, WebSphere, DB2

## Need 1 AIX Server

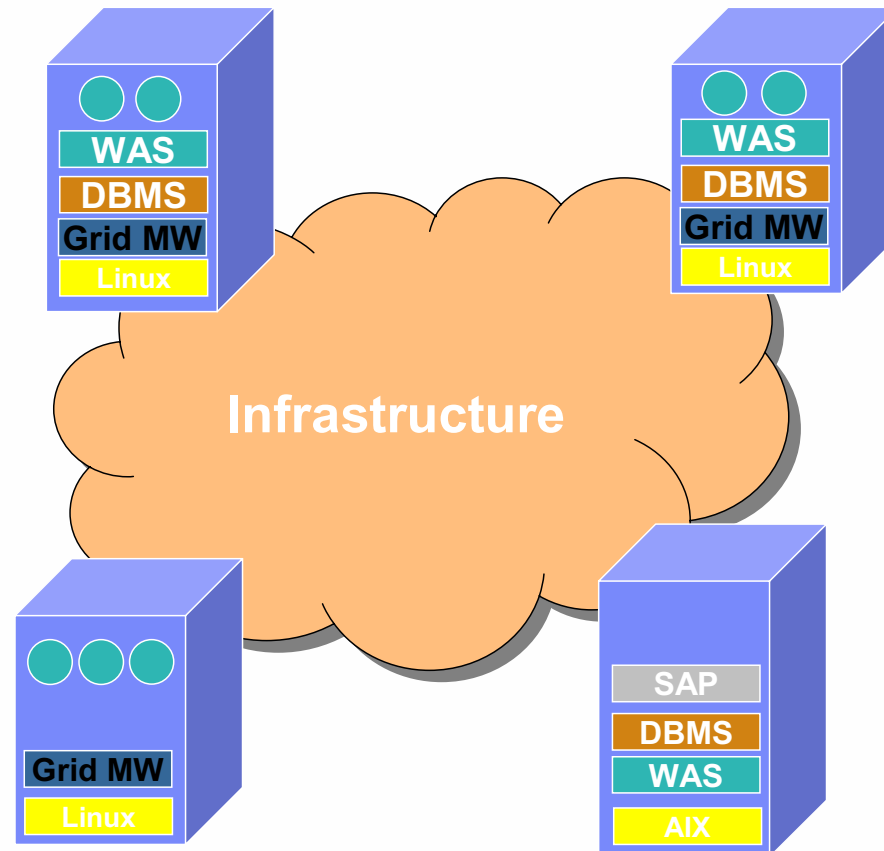
- § 32 Way Power4, 32 Gig of Memory
- § 1 Terabyte of Storage
- § AIX
- § SAP, DB2, WebSphere



# Orchestrated Provisioning



# Orchestrated Provisioning



# Orchestrated Provisioning

## § Senses conditions

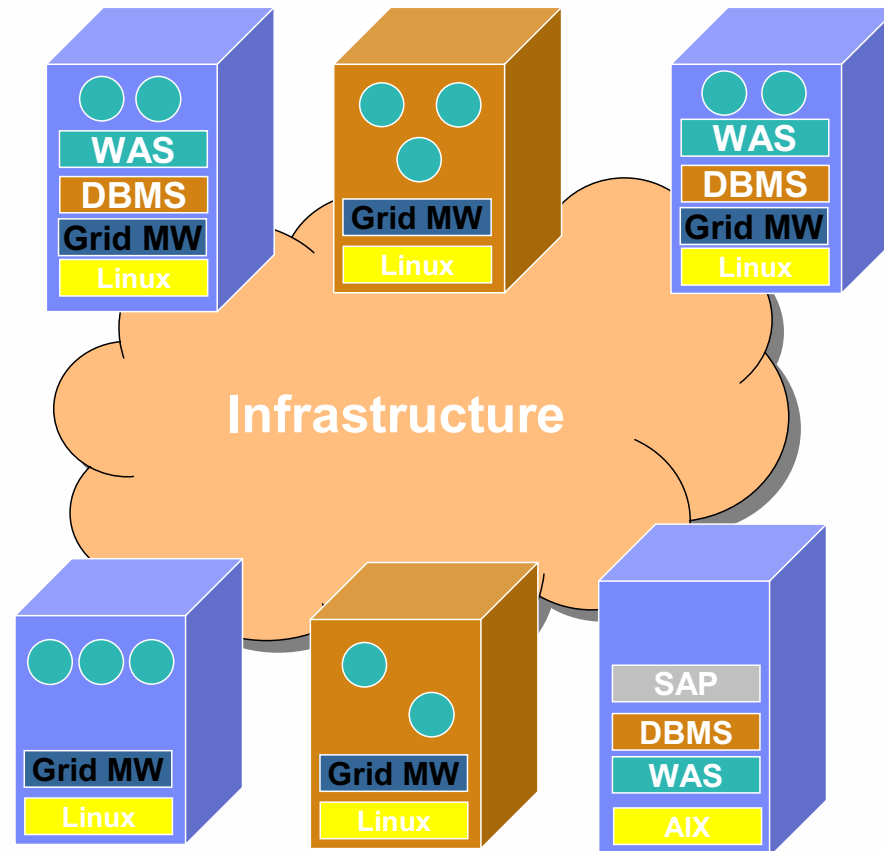
- The Grid needs more resources

## § Dynamically Responds

- **Provisioning** adds More resources to the Grid

## § Notifies Grid

- Jobs are scheduled onto the new resources



# Orchestrated Provisioning

## § Senses conditions

- The Grid needs more resources

## § Dynamically Responds

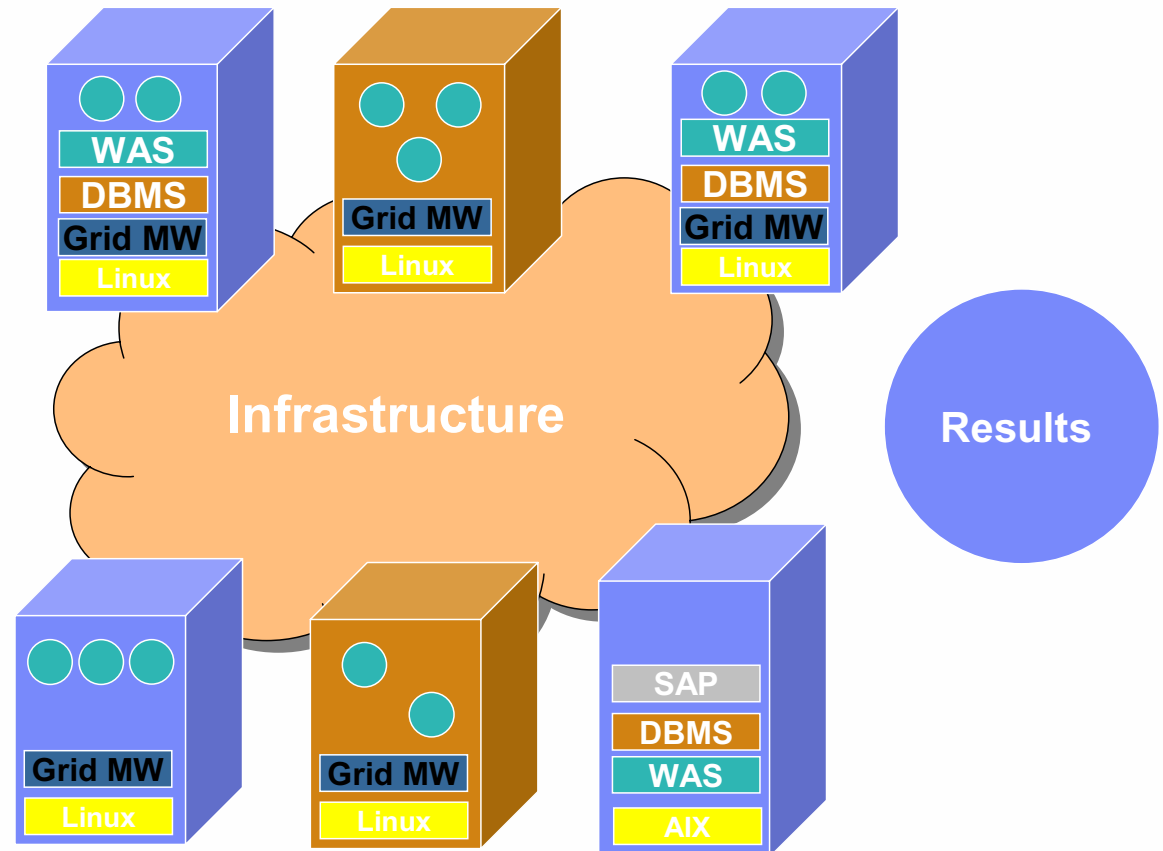
- **Provisioning** adds More resources to the Grid

## § Notifies Grid

- Jobs are scheduled onto the new resources

## § Job Completes

- Resources returned



# Which IBM products are key Grid components?

## § Open Grid Services Architecture

- **IBM Grid Toolbox for multi platform**

## § Storage

- **IBM SAN File System:** Virtualize heterogeneous SAN storage system and enable access from heterogeneous servers.

## § Systems Management & Policy Based Provisioning

- **IBM Tivoli Provisioning Manager V1.1:** Provisions and configures servers, operating systems, middleware, applications, and network devices.
- **IBM Tivoli Intelligent ThinkDynamic Orchestrator V1.1:** Sense, anticipates, plans, and controls responses to real-time production requirements.

## § Data Access

- **DB2 Information Integrator:** Designed to address customer requirements for integrating structured, semistructured and unstructured information effectively and efficiently across DB2, Oracle, Informix, Sybase, and MS SQL.

## § Job Management

- **Server Allocation for WebSphere Application Server (SAWAS):** Makes underutilized websphere Application Servers temporarily available for high performance parallel computing

## § Job Scheduling

- **LoadLeveler:** Used for dynamic workload scheduling
- **Tivoli Workload Scheduler for Virtualized Data Centers(ITWS –VDC):** Integrates with IBM LoadLeveler and Open Grid Services Architecture (OGSA) industry standard interfaces to provide the workload optimization and grid scheduling functions required by on-demand businesses

# IBM Grid Toolbox for Multiplatforms V3.0



§ A commercial derivative of the Globus Toolkit 3.0 with IBM Value Add

- A platform for the **Grid Developer** to develop and test grid service and grid applications.
- A platform for the **Grid Builder/Deployer** to host grid service and grid applications.

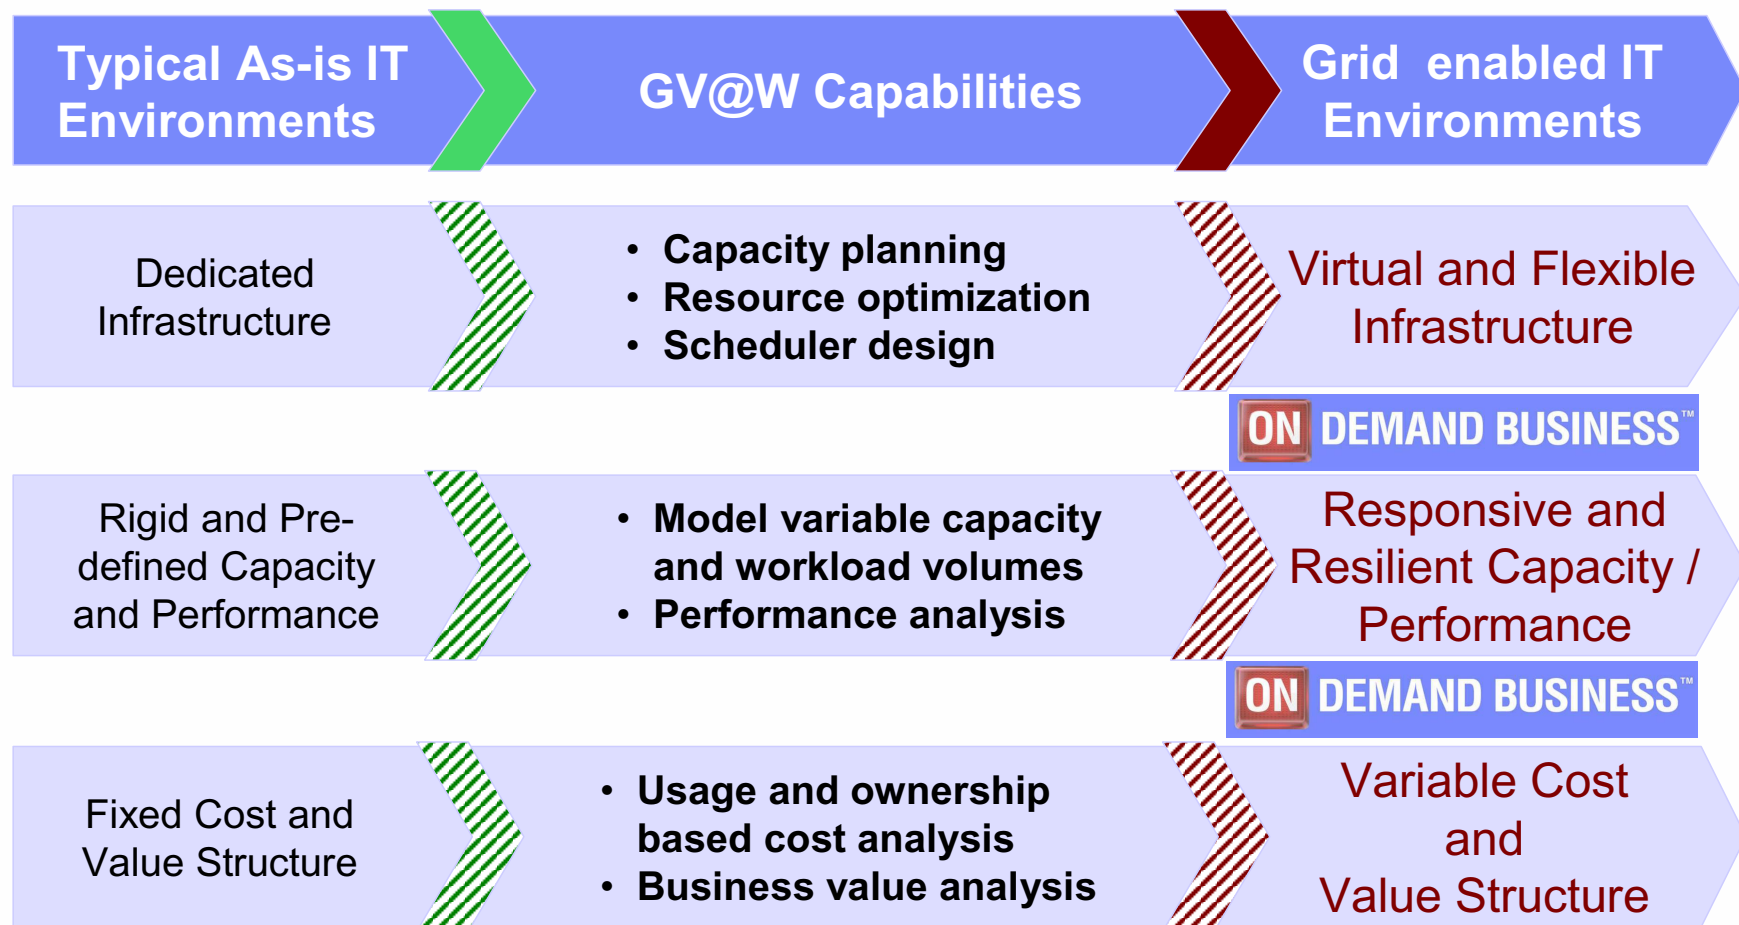
Customer scenarios used to design, document and test....

# IBM Grid Toolbox– the details....

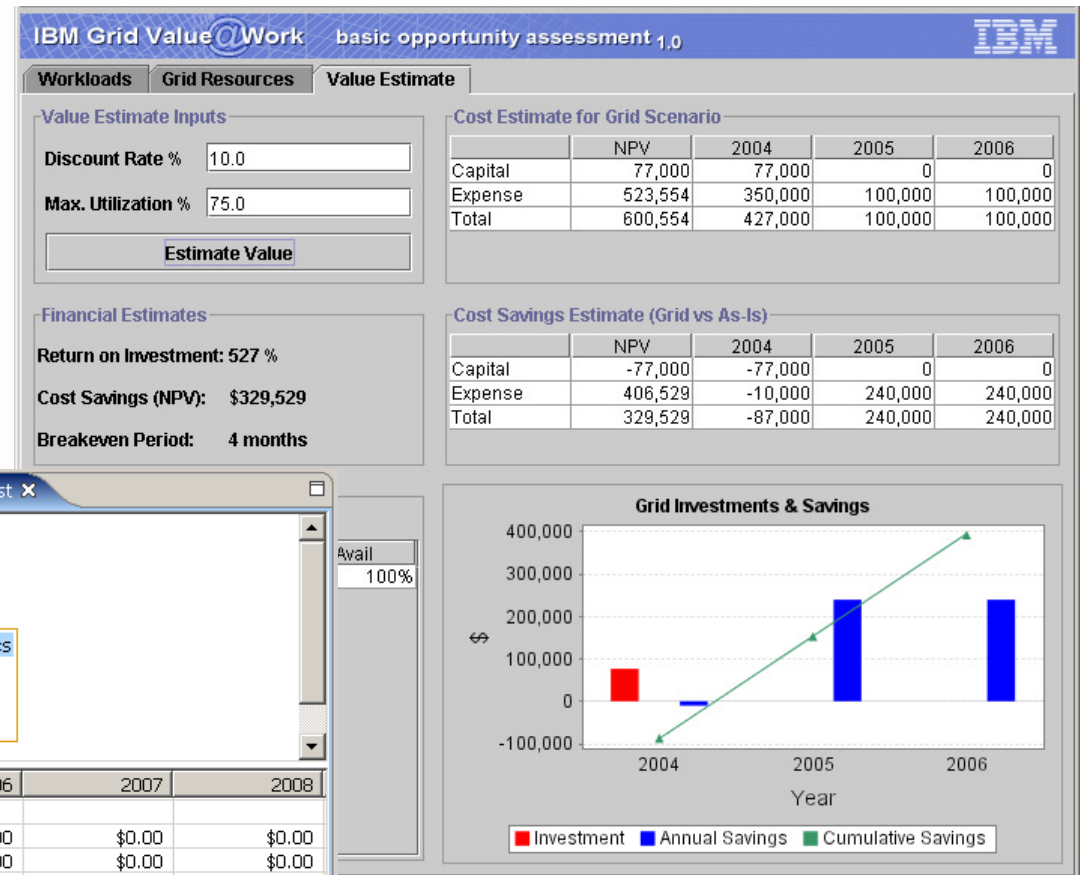
- § Integrated **wizard based installation**
- § A **grid service runtime** environment based on the GGF Grid Service (OGSI) specification for hosting grid services.
- § A web-based **management application** used to manage services within the runtime environment.
- § A digital certificate base grid security infrastructure (**GSI**).
- § Complete set of **configuration** and **administration** commands.
- § **Development tools** to build, package and use grid services.
- § Common (base) **grid services** for...
  - Discovery via Service Group
  - Policy Management
  - Common Management Models (CMM)
  - GT3.0 Program Execution Service ( GT3.0 GRAM - MMJFS, MJS, UHE)
  - GT3.0 Information Services (GT3.0 MDS - Index)
  - GT3.0 Data Management Services (Multi-RFT)
- § **Information Center** including tutorials to assist with the education and understanding of the technologies and capabilities packaged within the product.
  - Sample** grid services and applications demonstrating key capabilities.

# What is IBM Grid Value at Work

A Modeling and Financial Valuation tool for Grid Infrastructure Design



# The Grid Business Case



systest.gml Cost Compare: costmulti Cost: costmulti Business Case: systest

**systest**  
**Business Case: Grid Case**  
As-is Scenario: **Current** To-be Scenario: **Desktop Grid**

Inputs		Cash Flow Metrics		Accounting Income Metrics	
Discount Rate:	10%	ROI	1,091%	PTROC	255%
Cost of Capital:	8%	paybackMonths	15	ATROC	153%
Tax Rate:	40%	NPV	\$91,884	EVA	\$17,408
		IRR	360%		

Value Element	Present Value	2004	2005	2006	2007	2008
<b>Investments</b>	\$12,000.00					
capitalInv	\$12,000.00	\$12,000.00	\$0.00	\$0.00	\$0.00	\$0.00
capitalInvAvoided	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00
adjustedCapitalInv	\$12,000.00	\$12,000.00	\$0.00	\$0.00	\$0.00	\$0.00
<b>benefits</b>	\$103,884.15					
ITSavings	(\$36,509.59)	(\$27,000.00)	(\$3,000.00)	(\$3,000.00)	(\$3,000.00)	(\$3,000.00)
SGASavings	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00
COGSSavings	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00
RevenueImpact	\$140,393.74	\$30,750.00	\$32,287.50	\$33,901.87	\$35,596.97	\$37,376.82
totalBenefit	\$103,884.15	\$3,750.00	\$29,287.50	\$30,901.87	\$32,596.97	\$34,376.82
Depreciation		\$1,080.00	\$1,080.00	\$1,080.00	\$1,080.00	\$1,080.00
EBITImpact		\$2,670.00	\$28,207.50	\$29,821.87	\$31,516.97	\$33,296.82
<b>cashFlows</b>						
cashFlow	\$91,884.15	(\$8,250.00)	\$29,287.50	\$30,901.87	\$32,596.97	\$34,376.82
cumulativeCashFlow		(\$8,250.00)	\$21,037.50	\$51,939.37	\$84,536.34	\$118,913.16

Business Case

## Detailed Financial Analysis

## Basic Opportunity Assessment

# Grid Value at Work Use Cases

## Technology Adoption Lifecycle

**Awareness** ➔ **Interest** ➔ **Evaluate** ➔ **Trial** ➔ **Adopt**

Demonstrate financial value to senior execs using industry and application templates

**Financial Value Examples**

Identify suitable applications and resources for grid enablement

**Basic Opportunity Assessment**

Predict application performance on a grid design  
Show cost benefits

**Performance Simulation**

**Cost Analysis**

Develop business case for a grid implementation with cost and value assessment

**Cost Analysis**

**Business Value Analysis**

**Financial Analysis**

# What questions can be addressed by the tool?



I have a ***number of application*** workloads and IT systems that could be ***moved to a grid...***  
Give me a rough estimate of the investment and ROI



I am ***designing a grid***, show me the performance impact of:  
Alternate grid scheduling policies  
IT resource provisioning and orchestration policies  
Alternate grid configurations



Show how my ***grid design will support***:  
Random surges in workload volumes  
Unexpected capacity outages

more...

# What questions can be addressed by the tool?



Show me an **IT cost comparison** of:

- Current, various grid alternatives, non-grid upgrade options

- What are the detailed costs of building and operating a grid



Can you help me **assess the financial impact** of:

- Providing the capacity for increased workload volumes

- Providing faster response times to the IT end-user

- Increasing the frequency of running analyses



Develop a **business case** for a production grid environment:

- Return on capital invested, net present value, payback period and other financial metrics

- Accounting statement and cash-flow impact over time



Grid Computing

# Resources and 'Getting Started'



## Recommended Steps to Implement Grid Computing:

- 1. Education on grid computing and the benefits that it could bring to a department, agency, or enterprise**
- 2. Develop a 'vision' of the agency benefits that addresses BUSINESS issues..ie, enable better collaboration among scientists, sharing of information (and meta-information) that could enable greater productivity**
- 3. Executive sponsorship and buy-in is strongly supported**
- 4. Define and scope a specific "prototype" to get a successful proof-point that can be used to expand the infrastructure throughout the enterprise**
- 5. Conduct a grid workshop with the business and IT personnel to map out specifics**
- 6. Develop user- case scenarios that can be piloted and proven to show grid benefits**
- 7. Expand the prototype to include additional applications or departments**

# Organizational Span - Grid Gradations

---

## 1. One department or span of asset control

- No sharing with others
- No recovery from others
- Primary/secondary use
- Privacy, especially with desktop salvage

## 2. Two+ departments etc.

- Exploit others' resources
- Provide others resources
- WAN (maybe)
- As above

## 3. Two+ .com/.org's or utility

- Greater security, privacy, entitlement issues
- WAN
- Mutual audit, credentialing, certification,
- Contractual obligations to customers

# Grid Context

## Domain

- § Business Transformation Enablement
  - Time
  - Money
  - Content
  - Resiliency
- § Infrastructure
- § Utility
  - Consumer
  - Provider

## Grid Category

- § Intragrid
- § Extragrid
- § Intergrid

## Business Administration

- § Security
  - ID & Authenticate
  - Authorize
  - Entitle
  - Encrypt
- § Accounting
- § Billing

## Applications

- § Inhouse / 3<sup>rd</sup> party
- § Existing / New

## Resources

- § Compute-Memory-I/O
- § Data
- § Network
- § License
- § Sensors

## Equipment Source

- § Salvage / Acquisition
- § Utility / Overflow

## Equipment Type & OS

- § Workstation
- § Server
- § OS's : Win, CitWin, HP, zLin, ?xLin
- § Storage
  - Internal
  - External
  - Network

## Type

- § Compute
- § Data Content (storage & movement)
- § Data Centric (analysis & federation)
- § Deployment / Provisioning

## Workload Volumes

- § Timing
  - Well anticipated peaks,
  - Varying by application
- § Specifics
  - Requests & Sizes
  - Compute, Data Network

## Grid Programming Model:

1. Anywhere Batch
- \* Scheduled & Event Driven Workflow
2. Independent Batch
3. Subdividable Batch
4. Service
5. Parallel Service
6. Tightly Coupled HPC

## Data Granularity

- § Fine
  - Parameters
- § Medium
  - Database rows
- § Coarse
  - Files
    - Pre-positioned
    - On Demand
    - Shared

## Network Granularity

- § WAN vs LAN
- § Small
  - Request & Response
- § Medium
  - Request & Response
  - A little remote access
- § Large
  - Significant File & Database content
    - Not pre-positioned
    - Shared
    - On Demand

# If you let them, designers can design-out On Demand's benefits

---

Applications can best exploit capacity and resiliency in the On Demand world if:

- They have not been designed for inflexibility, non-scalability, nor to deadlock.
- Can exploit caching etc. if present
- They can run in a standard operating environment.
- Their owners are willing to share and give-to-get.
- So:  
Insist on quality design.  
Build on reference architectures.

It may take organizational and attitudinal change to get from  
**HERE** **to** **HERE.**

	Client / Server	J2EE	Grid
When is the target chosen?	Design Time	Deployment Time	Demand Time
For how long does the choice stick?	Forever	1-3 years	Minutes to hours
Who decides system settings and where to run?	Designer/ Developer	Operations	The system
Typically, how much sharing?	None	Some within assigned role e.g. network infrastructure, proxies, web servers. Not between production & test. Usually not database servers.)	Assignment as needed and prioritized

# Contra-indicators

---

## Business

- Unresolved audit, legal, security, privacy, or contractual problems
- There is another solution available with higher value and lower cost/risk/time required.
- Grid is appropriate technically, but Middleware + Application licensing is unaffordable

## Organization

- The other departments won't share.

## Network

- Network traffic is unacceptably high
- Data movement latency is unacceptably high

## Program:

- On an SMP, the existing program uses one instance with many threads that intensely update shared memory. Possibly reworkable with xSpaces, or MPI,
- The program exists, but it requires changes, you don't have the source code, and the vendor won't change it.



Grid Computing

# Grid Computing

## Grid Considerations and Questions



# Questions

- § Trends and state of the OSGI to WSRF transition within the grid community
- § Issues on grid implementations at other companies/agencies
- § Maturity of grid technology
- § Applications for grid in weather community
- § Clusters vs. Grids
- § On demand – higher order services
- § Surge capacity - DHS

After the daily morning homeland security briefing with the President, where credible intelligence was discussed that indicates the imminent release of a biological agent from airborne delivery mechanisms northwest of the San Francisco bay area, the national threat level is increased to **RED ALERT**.



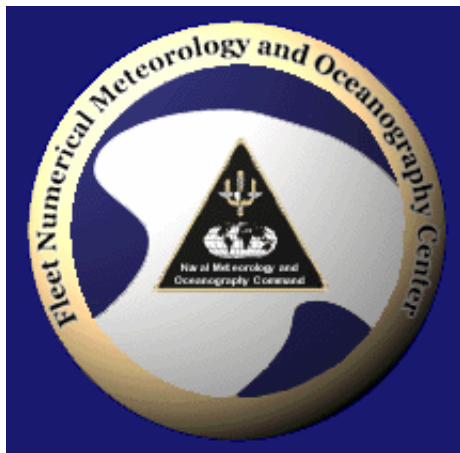
The airborne agents are to be gradually disbursed via hot air balloons, light civilian aircraft, remote control model airplanes, or short range missile delivery system. Weather conditions, wind speed, direction, humidity, and temperature are favorable for wide dispersal in populated areas.



A homeland security agent is responsible for providing information to law enforcement and local government regarding time to arrival at populated areas upon initial release detection, the direction of the agent cloud, and its altitude. This information will be used to schedule and mobilize emergency response and evacuation teams, provide more information to the public, and plan evacuation routes.



The agent logs into her homeland security portal. On the portal is a portlet for running dispersal simulations. The services invoked by the portlet are provided by FNMOC. FNMOC's modeling services Grid, remote operations Grid, and other data input products are used to execute the model.



Atmospheric Dispersion Model Portlet

Input Data Location:

GFTP://2.2.2.2/red-alert.xml

Resolution:

2.0

Results required in:

20 Minutes

Results Output Location:

GFTP://2.2.2.2/red-alert-out.xml

Input Data Transfer Bandwidth:

512Mbs

Status interval:

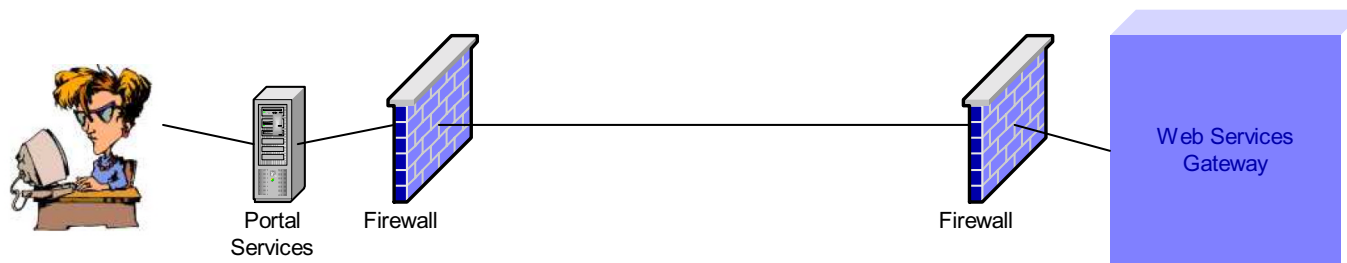
1 Minute

Run

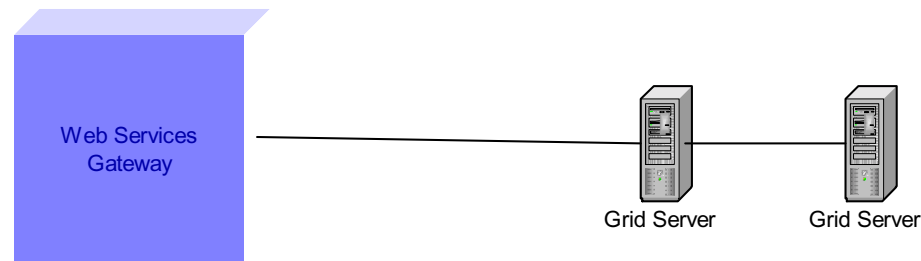


The agent invokes the portlet, which in turn invokes the modeling service at FNMOC:

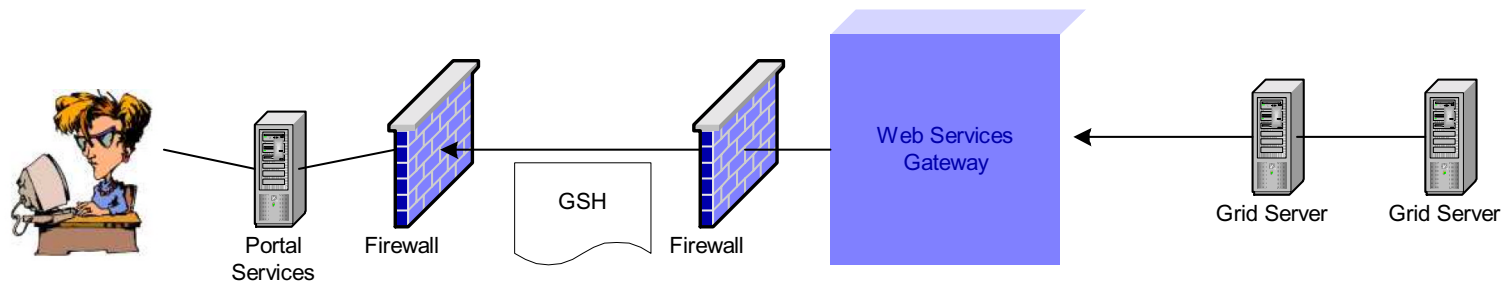
- The portlet generates a request for the modeling service in a SOAP message based on the WSDL definition exported by the gateway.



The gateway acts as a proxy and forwards the request to an available Grid server

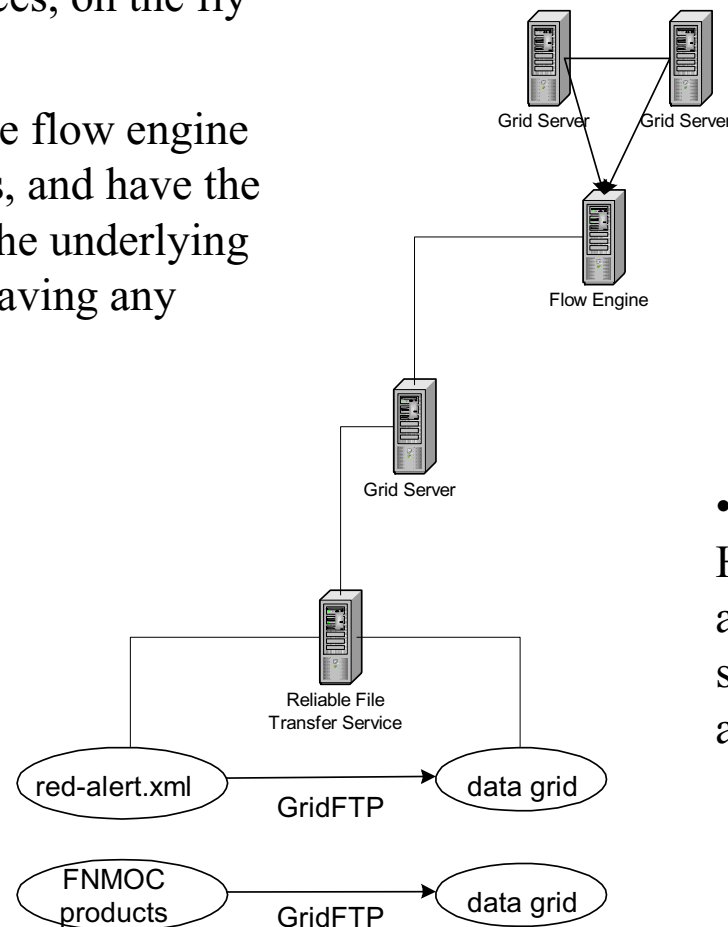


- The Grid servers invoke the Factory port type for the atmospheric model Grid service
- This instance of the service is added to the service registry of the Grid servers
- A Grid Service Handle (GSH) is created and sent back to the portal server/portlet instance
- The portlet uses the GSH to obtain a Grid Service Reference which is then used to subscribe to service data notifications
- Updated service data elements are sent back to the portlet at defined intervals



- The Grid service is a composed service in that it invokes another service, that of the flow for the model execution
- The flow engine in turn invokes other Grid Services.
  - Grid services, such as reliable file transfer service, are managed by other Grid servers
- Grid servers may dynamically deploy and execute the native service provider based on policy and resource availability
  - For example, to be able to move red-alert.xml along a 512Mbs circuit, it may be necessary for the Grid server to instantiate a Reliable File Transfer service, and possibly Grid FTP server instances, on the fly

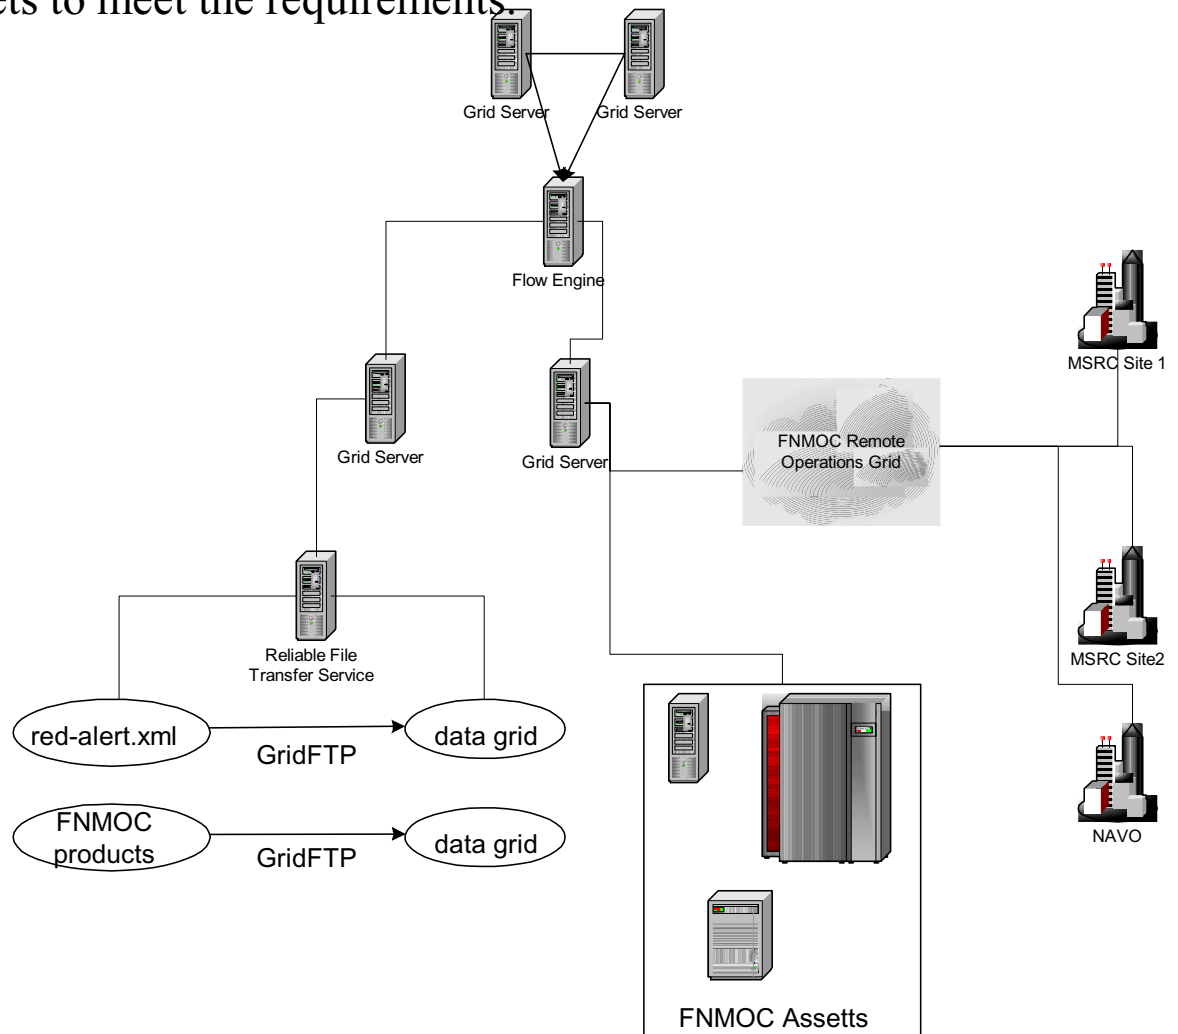
•Grid Services invoked by the flow engine pass along QoS requirements, and have the resiliency characteristics of the underlying Grid infrastructure without having any knowledge of the Grid.



•In this step of the flow, the HLS and FNMOC product files are sent to a data grid for staging and/or chunking.

- The flow engine then invokes the service to run the model
- The Grid server determines what resources are available, and what resources are busy
- Based on the resource requirements of the model, and the timeliness needed by the agent, and after checking policy to be sure the agent has high priority, the Grid server determines there are not enough available FNMOC assets to meet the requirements.

- The Grid server instructs all lower priority jobs running in engines to terminate and dispatches the model code to them
- The Grid server dispatches code to available engines
- After preempting processing at FNMOC, there are still not enough resources
- The Grid server performs engine allocation and preemption within the Remote Operations Grid
- Model code is deployed to the Remote Operations Grid
- Staged/chunked data is sent to engines





- When the service is completed, a final notification is sent to the portlet, letting the user know the simulation is complete.
- The various Grid service instances that were instantiated during the flow are destroyed.
- Preempted jobs are redeployed to engines in the Grid infrastructure.
- A large processing invoice based on resources used is sent to HLS J





Grid Computing

## ADDITIONAL INFORMATION and Resources

Nancy Brittle

[nbrittle@us.ibm.com](mailto:nbrittle@us.ibm.com)

[www.ibm.com/grid](http://www.ibm.com/grid)



# IBM Grid Computing Offerings for Public Sector

## § Government Grid offering:

- IBM Grid Offering for Information Access

## § Life Sciences Grid offerings:

- IBM Grid Offering for Analytics Acceleration
- IBM Grid Offering for Information Accessibility

## § Higher Education Grid offering:

- IBM Grid Offering for Research Collaboration

## § Grid Innovation Workshop

Grid Computing

IBM

IBM Grid Offering for Analytics Acceleration: Life Sciences

Envision new data volumes and requirements, managing IT resources and resources in distributed computing.

To facilitate bioinformatics and reduce modeling expenditures in the life sciences industry, view a range of tools, including open source code, in-house tools and applications from independent software vendors (ISVs), such as Jvarkit, ChARM, GPC, NGS, and HBS. These applications may need to run on a heterogeneous environment of PCs, workstations, clusters and high-performance computing systems. Grids often help harness high-throughput capacity, calculations and high-precision modeling capabilities.

Today, life sciences environment requires quick turnaround for large volumes of calculations and simulations. The IT environment handles multiple virtual, adaptable, flexible and scalable.

**Highlights**

- Cost reduction use of computing and storage resources across the enterprise
- Accelerated time to results and increased number of calculations processed
- Affordable, flexible IT resources meet for fast, large-scale analysis

**The challenges facing life sciences**

Genomic solutions in the life sciences are changing the way we live. These advances have the rapid ability to discover in genomics, proteomics and molecular biology that serve as the basis for individualized treatments. The advent of personalized medicine and the development of new drugs and treatments. Despite the complexity of pharmaceutical and biotech companies as well as healthcare organizations and academic research institutions, an increasingly widespread

IBM Grid Innovation Workshop

over the worldwide Internet or their own networks.

Although grid computing today is commonly used for the scientific and academic arena, commercial businesses stand to benefit from the new capabilities, increased collaboration and cost savings that can be achieved through grid technology.

**Introduction**

Grid computing is designed to be the solution that enables today's distributed computing resources into what is essentially a single, large virtual computer. By using a set of open standards and protocols, grids enable organizations to access data, processing power, storage and other heterogeneous computing resources.

IBM can help customers achieve their most challenging information problems by using grids to analyze and accelerate business opportunities. The scientific and academic communities recognize IBM as a leader in grid computing through its participation in The Global Project, the global

Executive Workshop	Executive Workshop For Work	Executive Workshop	Deliverables
<ul style="list-style-type: none"> <li>Define and provide an overview of grid computing concepts, evolution and available resources</li> <li>Provide a high-level overview of grid computing and its applications</li> <li>Explain the role of grid computing in the life sciences and its impact on the industry</li> <li>Identify and analyze key challenges in the life sciences grid computing</li> </ul>	<ul style="list-style-type: none"> <li>Describe detailed capabilities and how to use them to solve problems and achieve business goals</li> <li>Explain the role of grid computing in the life sciences and its impact on the industry</li> <li>Explain the role of grid computing in the life sciences and its impact on the industry</li> <li>Explain the role of grid computing in the life sciences and its impact on the industry</li> </ul>	<ul style="list-style-type: none"> <li>Define and provide an overview of grid computing concepts, evolution and available resources</li> <li>Provide a high-level overview of grid computing and its applications</li> <li>Explain the role of grid computing in the life sciences and its impact on the industry</li> <li>Identify and analyze key challenges in the life sciences grid computing</li> </ul>	<ul style="list-style-type: none"> <li>Detailed description of new grid computing capabilities and how to use them to solve problems and achieve business goals</li> <li>Explain the role of grid computing in the life sciences and its impact on the industry</li> <li>Explain the role of grid computing in the life sciences and its impact on the industry</li> <li>Explain the role of grid computing in the life sciences and its impact on the industry</li> </ul>

## IBM Grid Computing & OnDemand Design Centers

Provides an environment for customers to test their applications for:

- Grid Computing
- On-Demand computing
- Autonomic computing
- Information integration
- Web Services
- Linux

Grid Computing & OnDemand Design Centers are located around the world:

- USA - Poughkeepsie, NY
- USA - Silicon Valley, CA
- Europe - Montpellier, France
- AsiaPacific - Makuhari, Japan

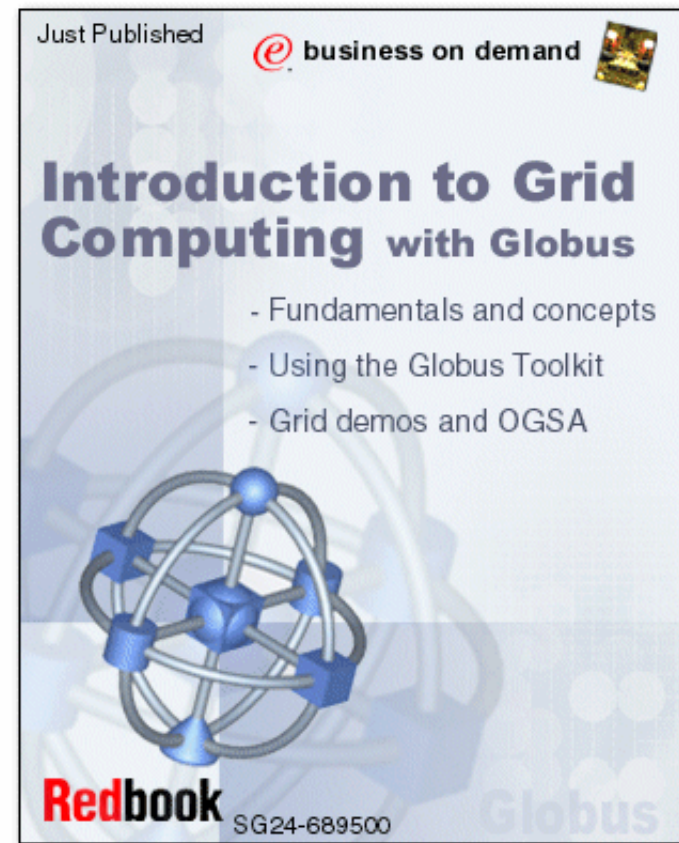
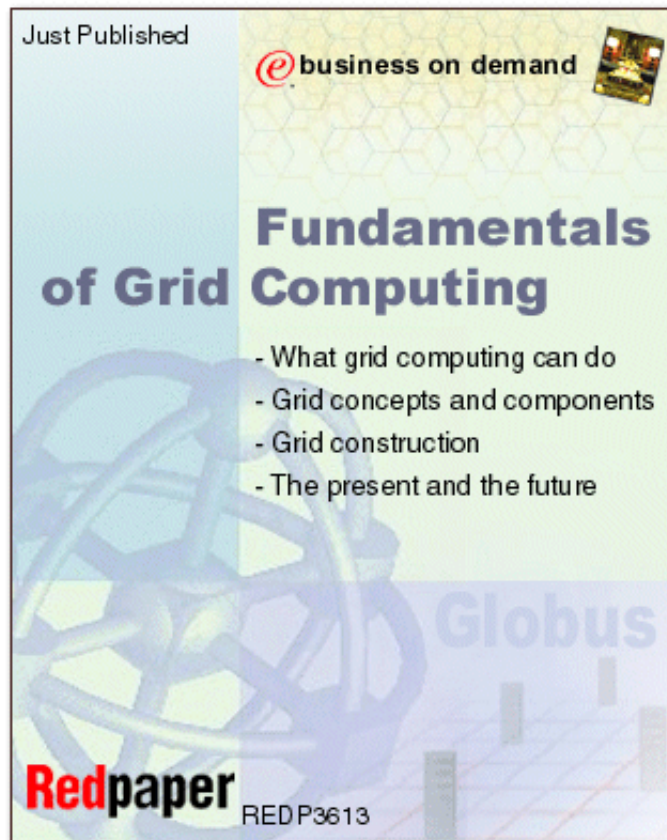


## **IBM Global Services can provide the following Grid Computing expertise**

- § **Overall Grid Computing Strategy**
- § **Identification and prioritization of business applications to “grid enable”**
- § **Calculation of Grid Computing Business Benefits**
  - ROI, TCO for the specific application area
- § **Selection of Grid Technologies**
  - Middleware recommendations, Selection of Standards and Platforms
- § **Grid Architecture Design**
  - Security, Management, Availability etc.
  - Operational considerations
- § **Grid Implementation – Pilot project**
- § **Grid Computing Education**
- § **Grid & Organizational impact assessment**

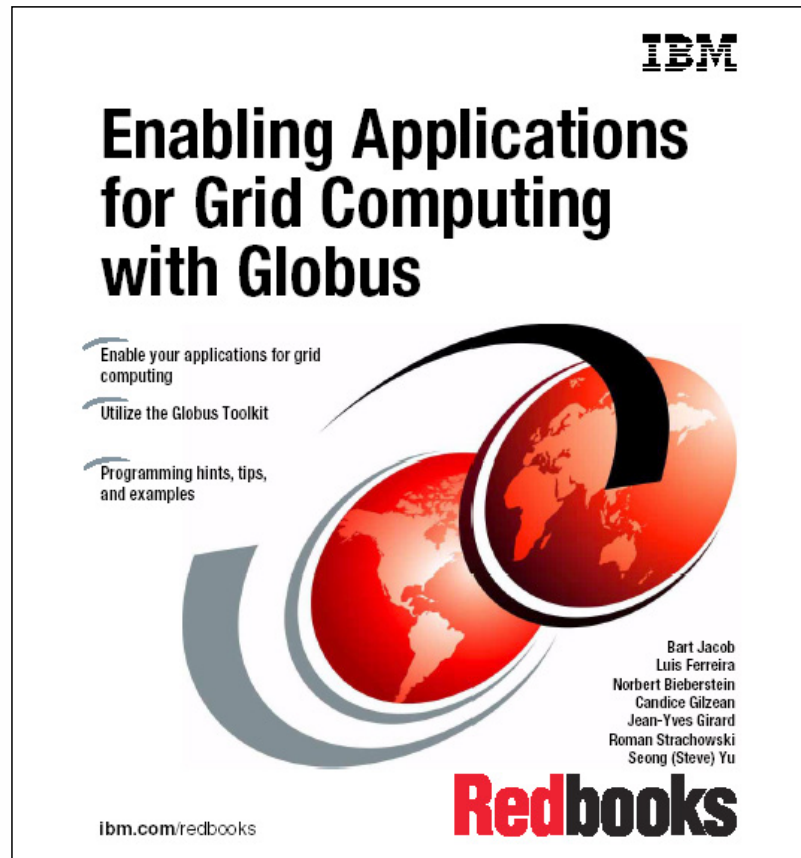


## Additional IBM Grid Information: Red Paper & Red Book



Download from [www.redbooks.ibm.com](http://www.redbooks.ibm.com)

# IBM RedBook: Grid Enabling Applications



Download from [www.redbooks.ibm.com](http://www.redbooks.ibm.com)

# IBM Grid Focus Areas

	Research & Development Grid	Engineering & Design Grid	Business Analytics Grid	Enterprise Optimization Grid	Government Development Grid
Description	Research intensive applications	Engineering and scientific applications to accelerate product design	Faster and more comprehensive business planning and analysis	Improve utilization, efficiency and business continuity	Support very large-scale IT infrastructures
Information Grid	Sharing of public data sources. Also supports use of shared compute resources.	Sharing design data across large multi-party projects.	Large scale data marts. Broadly distributed client data.	Virtualize distributed storage and data resources.	Provide large scale data sharing infrastructure for industry and scientific collaboration
Priority Sectors	Public, Industrial	Industrial	Financial Services, Public, Industrial	Financial Services, Public, Industrial	Public
Primary Buyers	<ul style="list-style-type: none"> <li>•Primary: LOB decision makers</li> <li>• Secondary: IT as an enabler</li> </ul>			<ul style="list-style-type: none"> <li>•Primary: IT decision makers</li> <li>•Primary: IT decision makers</li> </ul>	

Productivity =



= Efficiency



# European Aeronautic Defense and Space Company

## Challenge

- § EADS wanted to build an “on demand computing” model for the simulation tools used by their engineers to shorten analysis completion time and provide a single image of computer resources.

## Solution

- § Shorten the product design cycle with a Grid Computing platform based on:
  - § IBM <sup>^</sup>TM
  - § Linux
  - § Globus Toolkit
  - § GridXpert technology



### Technology Benefits:

- § More robust, scalable IT infrastructure that adjusts as requests fluctuate
- § Open standards permit easy integration of existing software

### Business Benefits:

- § Cut analysis and simulation time, while improving the quality of the output
- § Improve the productivity of the Design Office

# IN2P3

(Institut National de Physique Nucleaire et de Physique des Particules)

## Challenge

- The research institute needs to enhance the scalability, reliability and resilience of the existing grid environment to meet the large-scale, high-performance computing needs of new and existing users, as well as prepare for expansion to other grid environments throughout Europe.

## Solution

- IBM eServer™ pSeries® UNIX-based servers
- IBM eServer™ xSeries® Linux clusters
- Globus Tool Kit V3.0
- Storage capacity of up to thirty terabytes

*“We are extremely pleased about our collaboration with IBM. IBM’s technical expertise will allow us to rapidly achieve our goal to build a production-ready Grid to support our key research initiatives.”*

**-- Guy Wormser, Deputy Director - IN2P3**



### Benefits:

- Improved performance increases the number of compute intensive research projects.
- Enhanced environment increases the ability of the organization to contribute in key Life Sciences research.
- State-of-the-art, production-ready Grid allows European technological community to efficiently collaborate.

# eDiamond – UK Government and IBM Initiative

## Challenge

- § Build prototype of an eventual national database of:
  - § Mammography
  - § cancer
  - § Degenerative brain diseases.

## Solution

- § National access to archive of digital mammograms
- § Fewer patient files lost/misdirected
- § Allow immediate referral for expert consultation
- § Better prognosis
- § New avenues of prevention through correlation across demographics
- § Early identification and treatment of cancer
- § Improved healthcare service to patients

